

# Horowitz – Epistemic Akrasia

Brian T. Miller

November 15, 2017

## PREFACE TO THE HANDOUT:

Let's all agree that *endnotes* are an abomination. Please use *footnotes* in your essays that you hand in for this class.

[End preface]

## §1

Horowitz's goal in the paper: defend the *Non-akrasia Constraint*:

*Non-akrasia Constraint (NAC)*: it's never rational to be highly confident in:  $p$ , but my evidence doesn't support  $p$

Case:

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 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. (719)

What does Sam's total evidence support? Three plausible types of response:

	Acknowledges force of FOE?	Acknowledges force of HOE?	Respects NAC?
<b>Steadfast views</b>	yes	no	yes
<b>Conciliatory views</b>	no	yes	yes
<b>Level-splitting views</b>	yes	yes	no

## §2 – Level-splitting views

The literature contains four ways to motivate level-splitting:

### §2.1 – Williamson

Williamson's case (described by SH):

Long Deduction: Suppose a rational agent comes to know a series of claims and competently deduces their conjunction, *C*. On Williamson's view, she can come to know *C* by these means. But suppose further that the rational agent also knows that, since most people's memories and logical abilities are limited, people in her situation often make inferential errors while completing long deductions. It can then be highly probable on her evidence that she has made such an error, and thus that she does not know the conjunction. Nevertheless, since she has in fact competently deduced *C*, Williamson holds that she does know *C*. In a case like this, it will be certain on her evidence that *C* is true—for Williamson, knowledge of *C* requires assigning *C* evidential probability 1—but also highly probable on her evidence that she does not know *C*. So she should be highly confident in *C*, **and highly confident that she does not know that *C***. (720, emphasis added for later reference)

SH fleshes out the story in order to turn the boldface part into an explicit violation of the *Non-akrasia Constraint*, i.e. to turn *C*, and highly confident that she does not know that *C* into *C*, and my evidence doesn't support *C*

Problem:

- in order to know *C*, one must be rationally certain that it's true (according to Williamson)
- doubt about whether the evidence supports *C* comes from suspicion that one has made an inferential mistake
- inferential mistakes usually lead to false conclusions,
- so, this case won't lead you to *both* be certain that *C* and doubt that your evidence supports *C*

Proposal:

- as a matter of fact, *S* believes *C* on the basis of *E*<sub>1</sub>, *S* is certain that *E*<sub>1</sub>, and *E*<sub>1</sub> entails *C*
  - so, *S* should be rationally certain that *C*

- but S has strong evidence that her evidence for C is E<sub>2</sub>, and E<sub>2</sub> does not entail C<sup>1</sup>
  - So, S should doubt that her evidence supports C

<sup>1</sup> TW think's it's possible to mistake your evidence, and to mistake what your evidence supports.

So, S should believe: C and my evidence doesn't support C

So, TW rejects the *Non-akrasia Constraint*.

## §2.2 – Lasonen-Aarnio

Case:

**Mental maths:** My friend and I have often amused ourselves by solving little math problems in our heads, and comparing our answers. We have strikingly similar track records: we are both very reliable at doing mental maths, and neither is more reliable than the other. We now engage in this pastime, and I come up with an answer to a problem, 457. I then learn that my friend came up with a different answer, 459. (315)

S's FOE: her calculation that the answer is 457.

So's HOE: the fact that her friend didn't get 457, which indicates the S calculated incorrectly, and hence that her evidence doesn't support her conclusion.

It's tempting to think that S's HOE defeats her justification for believing that the answer is 457.

That tempting thought is what Lasonen-Aarnio called the thesis of *Higher-order Defeat*.

But that leads to a puzzle with no satisfactory response (as we saw last week)

So, we should reject *Higher-order Defeat*: the HOE doesn't defeat the belief based on the FOE

So, in this case S should believe: the answer is 457, but my evidence doesn't support my believe that the answer is 457

Because we've rejected *Higher-order Defeat*, the latter conjunct doesn't defeat the former (or vice versa) – it's perfectly rational to believe both at the same time.

So, Lasonen-Aarnio rejects the *Non-akrasia Constraint*.

## §2.3 – Wedgewood

Beliefs formed via inferences manifesting a 'necessarily rational disposition' are themselves justified.

[I]t is the real nature of this internal process—the fact that it is a process of competent inference—and not the higher-order beliefs that the thinker has, or even the beliefs that the thinker is justified in having, about the nature of that process, that is crucial to the rationality of the mental event that results from that process. (723)

So, if Sam in Sleepy Detective infers that the thief was Lucy, and that inference manifests a necessarily rational disposition, then he is justified in believing that the thief was Lucy.

It may also be the case that Sam is justified in believing that his evidence does not support that conclusion, since Alex called his attention to his poor record of reasoning when tired.

So, Sam is justified in believing:  $p$ , and my evidence does not support  $p$ .

So, Wedgewood rejects the *Non-akrasia Constraint*.

#### §2.4 – Weatherson and Coates

Both motivated by a kind of *normative externalism*: what evidence supports is independent of what you *think* your evidence supports.

[T]here are facts about which hypotheses are supported by which pieces of evidence, and ...rational agents do well when they respond to these epistemic facts. ...[T]hese facts retain their normative significance even if the agent has reason to believe that she's made a mistake in following them. That is, if an agent's judgment conforms to the correct norms of judgment, then even if she has evidence that she is not good at judging, she should stick to her judgment. (724)

So, having evidence that supports  $p$  is consistent with having *misleading* evidence that your evidence does not support  $p$ .

So, in that case you should believe:  $p$ , but my evidence does not support  $p$ .

So, Weatherson and Coates reject the *Non-akratic Constraint*.

#### §2.5 – Extreme Level-splitting

The authors discussed are committed at least *some* degree of level-splitting: the rational level of confidence in  $p$  comes apart from the rational level of confidence you are highly confident you should have.

In extreme cases, they come apart a lot: your rational attitude is one of high confidence in  $p$ , high confidence that your evidence does not support  $p$

But the considerations adduced in favor of moderate level-splitting views also support extreme level-splitting views.

So, SH attacks extreme versions, thereby attacking the motivations for the weaker versions.

### §3 – Immediate problems

#### First problem:

According to Level-splitters, Sam should believe:

- (i) Lucy is the thief
- (ii) My evidence doesn't support: Lucy is the thief

But that sounds wrong:

- Upon reflection, what is Sam supposed to think about his belief that Lucy is the thief? 'My evidence doesn't support that belief, beliefs unsupported by evidence tend to be false, in this case I just got lucky and came to believe truly'? That sounds obviously incorrect.
- Imagine Sam receives the HOE first: he believes he's about to assess some FOE unreliably. Then he evaluates the FOE, and comes to believe that Lucy is the thief on that basis:

According to the Level-Splitting view, he should be highly confident that P. "I thought I was going to judge falsely," Sam might say to himself, "but I must have I lucked out! I judged that P, and P is true." (726)

That sounds obviously incorrect.

#### Second problem:

If the level-splitters are right, then Sam is in a position to infer that his FOE is misleading in a problematic way.

In the Sleepy Detective case,

- in the absence of FOE, Sam's confidence in p (=that Lucy is the thief) should be *low*
- with the FOE, his confidence in p should be *high*
  - Sam concludes: p is true
- by the HOE, Sam should believe that the FOE doesn't support p

- Sam reasons:  $p$  is true. But all of my FOE relevant to  $p$  does not support it. It supports low confidence in a true proposition,  $p$ , and therefore high confidence in a false proposition,  $\neg p$ . So, my FOE is misleading.

Case to illustrate:

Sam is driving in to work one morning, and plans to evaluate some evidence,  $E$ , when he gets there. He hasn't seen  $E$  yet, but as he is driving, he rationally thinks to himself: "Whatever  $E$  is, it is most likely not misleading—after all, good evidence usually supports the truth." When he walks in the door, Alex tells him (falsely) that she has spiked his coffee with a reason-distorting serum. "That's too bad," Sam thinks. "Now I won't be able to evaluate my evidence correctly; I'm unlikely to find out the truth about the jewel thief." Then Sam opens the door to the evidence room, and sees  $E$  before him.  $E$ , in fact, supports  $P$ . Just as in Sleepy Detective, Sam evaluates  $E$  rationally, concludes that  $P$ , and takes his higher-order evidence to indicate that  $E$  does not support  $P$ . "Aha!" he exclaims, " $P$  is true! But this evidence doesn't support  $P$ ; it's misleading. That's too bad." Then Sam thinks another minute, and smiles. "Actually, it's a good thing Alex spiked my coffee. If I had evaluated my evidence correctly, I would have ended up with a false belief." (726)

This is an illicit way to decide that your evidence is misleading, and level-splitters are committed to it.

### Third problem:

If Sam's beliefs are what the level-splitters say they should be, then he might act as follows:

*Sam:* I'd bet that it's Lucy. I'll give you 9:1 odds.

*Alex:* But you were so sleepy when you were working last night!  
How can you be so sure that the evidence supports her guilt?

*Sam:* Oh, I'm not. Since you told me that I'm so bad at evaluating evidence when I'm tired, I doubt the evidence supports Lucy's guilt much at all. If I were to bet on what the evidence supported, I might give you 1:9 odds that it's Lucy, but certainly not 9:1.

*Alex:* So why are you offering 9:1 odds?

*Sam:* Well, I really shouldn't be offering such strong odds. I shouldn't be so confident that she's the thief: the evidence isn't in my favor. But on the other hand, she is the thief! That's what we're betting on, right?

That's a crazy way to act, but it follows naturally from the beliefs recommended by the level-splitter.

What's more:

...if Sam can act on his akratic beliefs, his odd-looking behavior won't be restricted to betting. Should Sam tell newspaper reporters that Lucy is the thief? Recommend that the cops raid her apartment? Send her to jail? If he is rationally highly confident that she is guilty, it seems that all of these actions could be warranted. But if Sam is asked to justify or explain his behavior, he will be at a loss: as far as he thinks, his evidence doesn't support any of it. In fact, he might think that he shouldn't be acting in these ways. (728)

Seems problematic.

#### §4 – *Problems with extending level-splitting*

General question for level-splitters: when, if ever, do the levels interact?

In this section SH argues that there is no unproblematic answer for level-splitters.

##### §4.1 – *No interaction ever*

SH: this position is a non-starter.

Here's a case in which some interaction seems obviously appropriate:

*Alex:* I've figured it out. All of the evidence points to Veronica's guilt!

*Sam:* Well, I don't care what the evidence points to. Let me see the evidence, so I can figure out who did it.

*Alex:* Veronica did it.

*Sam:* You could have said so in the first place! I thought you were just telling me what the evidence supported.

So, at least *sometimes* evidence interacts across levels

##### §4.2 – *Interaction only in the absence of relevant first-order evidence*

The above case illustrates that HOE and FOE interact at least *sometimes*. But when?

Proposal:

*Proxy*: HOE is relevant to first order beliefs only insofar as it serves as a proxy for first-order beliefs

When you have no FOE but you learn the HOE fact that the evidence supports *p*, then the HOE gives you some reason to believe *p*

But, when you have the FOE for *p*, and then you obtain HOE that the FOE doesn't support *p*, then the HOE has no rational effect: it's *screened off* by the FOE

This approach seems to get the right result in the case above.

Problem: Proxy seems to be missing something:

*Case 1*: Sam is working to find a serial killer. He has his evidence, *E*, set out in front of him, and he has carefully looked through all of it. *E* supports the proposition *M*: that Marie is the killer. Sam has not yet finished his analysis of *E* and has not reached a conclusion about what *E* supports. Nevertheless, Sam just happens to believe—for no particular reason—that Marie is the killer.

*Case 2*: Sam and Alex are working together to find a serial killer. They have their evidence, *E*, set out in front of them, and they have both carefully looked through all of it. *E* supports the proposition *M*: that Marie is the killer. Sam has not yet finished his analysis of *E* and has not reached a conclusion about what *E* supports. But Alex finishes her analysis first, and exclaims, "I've got it! The evidence points to Marie." On the basis of Alex's testimony, Sam comes to believe that Marie is the killer. (730-1)

Judgment: Sam's belief is *propositionally* justified in both cases, but only *doxastically* justified in case 2

Why the difference? Presumably because Sam has HOE in case 2 that he lacks in case 1

BUT: in both cases he has the same FOE, and according to Proxy, anytime one has the relevant FOE, the HOE is screened off and irrelevant

So, the asymmetry between case 1 and case 2 can't be explained by appeal to Proxy

So, the Proxy-based account is missing something

Proposed addition to Proxy: HOE also affects the rationality of FO belief by bearing upon the support relation between the FOE and the belief.

Problem: although this produces the right answer in case 2, it produces the wrong answer (from the level-splitter's point of view) in



Sleepy detective

More generally, seems like this type of interaction is precisely what the level-splitters want to reject.

#### §4.3 – *Interaction in the long run*

According to level-splitters, in Sleepy Detective Sam should end up believing:  $p$ , and my evidence doesn't support  $p$

Suppose Sam found himself in this situation over and over again: the according to level-splitters he should believe:  $p$ ,  $q$ ,  $r$ , ..., and my evidence doesn't support  $p$ , my evidence doesn't support  $q$ , my evidence doesn't support  $r$ , ...

Upon reflecting, what should Sam think about this?

First option: he should remain akratic

1. he's extremely lucky
2. he's got a guarding angel
3. he's psychic
4. his beliefs were supported by his evidence after all

Problem: (4) is the best explanation, but the Remain-Akratic level-splitter can't say that.

Second option: he shouldn't remain akratic – he should become more confident that he beliefs line up with what his FOE supports

Problem: that kind of inference is a case of bootstrapping: assessing the reliability of a belief process by using the process itself, without employing any independent checks.<sup>2</sup>

#### §5 – *Bridge principles, reasoning, and action*

The arguments in §3 and §4 depend on the following bridge principles:

*Reasoning Bridge Principle:* It is rational to reason from any proposition,  $P$ , just in case one is rationally highly confident in  $P$

*Action Bridge Principle:* It is rational to act on any proposition,  $P$ , just in case acting on  $P$  maximizes expected value. (734)

<sup>2</sup> Note, however, that this problem doesn't arise because of a conflict between FO beliefs and HO beliefs – you'd get the same problem with someone who forms FO beliefs while lacking HO beliefs. The only way out of the pickle is to adopt a kind of neo-rationalism, on which you actually have justified beliefs in all the HO propositions, which is dubious.

Level-splitters could respond to the argument by rejecting the principles.

Problems with that approach:

1. Both principles are highly plausible, and they're both consistent with level-splitting, so rejecting them is an additional problematic commitment for the level-splitter
2. Even if the bridge principles are false, you still get a problematic result:
  - Suppose Sam rationally believes:  $p$ , and my evidence doesn't support  $p$  (as the level-splitter thinks possible)
  - If the Reasoning Bridge Principle is true, then Sam is rational in inferring that his evidence for  $p$  is misleading (this was the 'second problem' from §3)
  - But even if the Reasoning Bridge Principle is false, Sam's beliefs still *entail* that his evidence of  $p$  is misleading
  - So, if the level-splitter wants to avoid saying that Sam is in a position to rationally believe that his evidence for  $p$  is misleading, she must say that Sam is not in a position to rationally believe that logical consequences of his rational beliefs, a strange kind of closure-failure

### §6 – Possible exceptions?

So epistemic akrasia is always irrational, right?

Not so fast: on the assumption that one can be rationally uncertain about what one's evidence is, then it's sometimes rational to be akratic.

Case:

**Dartboard:** You have a large, blank dartboard. When you throw a dart at the board, it can only land at grid points, which are spaced one inch apart along the horizontal and vertical axes. (It can only land at grid points because the dartboard is magnetic, and it's only magnetized at those points.) Although you are pretty good at picking out where the dart has landed, you are rationally highly confident that your discrimination is not perfect: in particular, you are confident that when you judge where the dart has landed, you might mistake its position for one of the points an inch away (i.e. directly above, below, to the left, or to the right). You are also confident that, wherever the dart lands, you will know that it has not landed at any point farther away than one of those four. You throw a dart, and it lands on a point somewhere close to the middle of the board. (736)

Suppose the dart lands at  $\langle 3,3 \rangle$ . What should you believe about where it landed?

Plausibly, you shouldn't be *certain* that it landed at  $\langle 3,3 \rangle$ , but you can narrow down the possibilities to that point on the ring of points around it.

Now consider the proposition Ring: that the dart landed on one of  $\langle 3,2 \rangle$ ,  $\langle 2,3 \rangle$ ,  $\langle 4,3 \rangle$ , or  $\langle 3,4 \rangle$ . (That is, that it landed on one of the points in the ring around  $\langle 3,3 \rangle$ .) Your credence in Ring should be .8. But that level of confidence in Ring is only rational if the dart actually landed on  $\langle 3,3 \rangle$ . If the dart had landed on some other point, the rational credence distribution would be centered on that point instead of  $\langle 3,3 \rangle$ , and your rational credence in Ring would be lower than .8. In particular, if the dart landed at any of  $\langle 3,2 \rangle$ ,  $\langle 2,3 \rangle$ ,  $\langle 4,3 \rangle$ , or  $\langle 3,4 \rangle$ , it would be rational for you to have .2 credence in Ring. Suppose you have reflected on your situation, and you know what would be rational to believe in these various evidential situations. So you should be .8 confident of Ring, but also .8 confident that your credence in Ring should be .2. This means you should be epistemically akratic: you should be highly confident of both "Ring" and "my evidence supports  $\neg$ Ring". (737)

Suppose that's right, and it's sometimes rational to be akratic. Does that mean that Sam should be akratic in Sleepy Detective?

Important difference between Dartboard and Sleepy Detective:

- In Sleepy Detective, as in most cases, we expect the evidence to be truth guiding: generally, when the evidence indicates high rational confidence in  $p$ ,  $p$  is true
- But Sam can infer that his FOE is misleading, i.e. *not* truth guiding:

Epistemic akrasia seems so odd in this case, in part, because it involves ignoring the reasonable background expectation that one's evidence is truth-guiding. (738)

- Dartboard is a weird case precisely because the evidence is *not* truth-guiding wrt propositions such as Ring. It's *falsity* guiding: when the evidence supports high confidence in Ring, Ring is probably false ( $p(\text{ring})$  is high iff the dart lands at  $\langle 3,3 \rangle$ , in which case Ring is false).

You should think that you should have low confidence in Ring precisely because you should think Ring is probably true—and because your evidence is falsity-guiding with respect to Ring. Epistemic akrasia is rational precisely because we should take into account background expectations about whether the evidence is likely to be truth-guiding or falsity-guiding. (738)

- We get a different answer in Sleepy Detective:

...suppose that Sam starts off rationally expecting his first-order evidence to be truth-guiding. Suppose he also has HOE (Alex's testimony) suggesting that his likelihood of interpreting the FOE correctly is no better than chance. Even after seeing the FOE, then, Sam should remain highly uncertain about what it supports. And because he still expects his evidence to be truth-guiding, his uncertainty about what the evidence supports should bring with it uncertainty about which suspect is guilty. So even after seeing the evidence, Sam should be highly uncertain about which suspect is guilty. For any particular suspect, including Lucy, he should end up with low credence that that suspect is guilty. And this is exactly what the Conciliatory view recommends. (739)<sup>3</sup>

### *Conclusion*

Due to cases like Dartboard, it's sometimes rational to believe 'p and my evidence doesn't support p'. So, the non-akratic condition is false.

But, counterexamples to the non-akratic condition are strictly limited to cases in which the evidence is falsity-guiding: like Dartboard, and unlike Sleepy Detective

<sup>3</sup> Does this require that we actually have all beliefs about what the evidence supports and about whether we are interpreting the evidence correctly? Why all the talk of 'expecting' the evidence to be truth-guiding, rather than *believing* that the evidence is truth-guiding? Looming regress? Lewis Carroll problems? Psychological implausibility? Harmless linguistic variation?