

Can Worsnip's Strategy Solve the Puzzle of Misleading Higher-Order *Apparent* Evidence?

March 8, 2018

Paul Silva Jr.
psilvajr@gmail.com
forthcoming in Inquiry

Abstract

It's plausible to think that we're rationally required to follow our total evidence. It's also plausible to think that there are coherence requirements on rationality. It's also plausible to think that higher-order evidence can be misleading. Several epistemologists have recognized the puzzle these claims generate, and the puzzle seems to have only startling and unattractive solutions that involve the rejection of intuitive principles. Yet Alex Worsnip (Philosophy and Phenomenological Research, forthcoming) has recently argued that this puzzle has a tidy, attractive, and independently motivated solution that involves rejecting the claim that we're rationally required to follow our total evidence. In what follows I argue that this solution fails to solve the fundamental problem for rationality.

Acknowledgements: I am grateful to the Inquiry referees for valuable feedback that greatly improved this paper.

1 Evidence and Coherence: the Apparent Conflict

Many epistemologists have regarded it as a platitude of epistemic rationality that it requires one's doxastic attitude towards a proposition—henceforth 'D(p)'—to track one's total evidence.¹ Here's Worsnip's (forthcoming) way of capturing this thought:

EVIDENTIAL REQUIREMENT (ER). If S's total evidence supports D(p), then rationality requires of S that she takes D(p).²

(ER) leads only to a very weak version of evidentialism. This is for two reasons. First, it leaves room for rational beliefs that are not based on evidence. So evidence is not necessary for rational belief. Second, (ER) is neutral on the nature of evidence. So it's not only internalist evidentialists who have a stake in this evidentialist principle, but anyone who thinks our attitudes ought to track our evidence. For both these reasons (ER) is consistent with a wide range of positions in epistemology.

(ER) is an extremely intuitive thesis. Suppose you study the history of Europe and *all* your evidence supports believing that Napoleon won many battles but lost a war. It is rational for you to believe that, and not believing that would be irrational. It's hard to imagine cases where one's total evidence supports believing p though one can rationally not believe p. Even if an epistemologist insisted that there are *some* cases where it's rational to resist one's total evidence, they should still insist that it's *not always* rational to do so. And that's enough to get the puzzle going.

¹'D(p)' is taken to range only over the coarse-grained attitudes of belief, disbelief, and suspension, and no uniqueness thesis is assumed. As Worsnip (forthcoming) points out, this way of setting things up allows us to avoid unnecessary and distracting technicalities in the assessment of the core philosophical problem to follow. See section I of his paper for further caveats and cautionary notes.

²For defense of this sort of principle see Conee and Feldman (2004) and Lasonen-Aarnio (forthcoming).

Many epistemologists have maintained that rationality requires one to be coherent in at least the following sense:

INTER-LEVEL COHERENCE (ILC). Rationality requires that:

(i) S believes that her total evidence supports D(p) only if she takes D(p); and

(ii) S believes that her total evidence does not support D(p) only if she does not take D(p).³

On the surface (ER) and (ILC) seem to gel quite nicely, expressing two distinct yet plausible principles concerning the demands of epistemic rationality. (ER) tells one to follow her evidence; (ILC) tell one to avoid akratic pairs of attitudes. Yet the puzzle surfaces with the possibility of misleading higher-order evidence. Take *higher-order evidence* to be evidence about what one's evidence supports. Like virtually all evidence, it seems possible that higher-order evidence can be *misleading*. That is:

POSSIBILITY OF ITERATIVE FAILURE (PIF). It is possible that:

(i) S's total evidence supports D(p); and

(ii) S's total evidence supports believing that her total evidence does not support D(p).⁴

I will give an intuitive case in which (PIF)(i) and (ii) obtain below. But apart from that I will not be defending (PIF). I have nothing to add to existing defenses of (PIF)–or (ILC) or (ER) for that matter. My primary aim is to demonstrate that rejecting (ER) will not resolve the central problem.

³See Worsnip (forthcoming) for a recent defense of this. For others who support (ILC) or something much like it see Elga (2005), Feldman (2005), Christensen (2007, 2013), Broome (2013, 98), Greco (2014), Horowitz (2014), and Titelbaum (2015).

⁴For defense of this see Worsnip (forthcoming), Lasonen-Aarnio (forthcoming), Littlejohn (forthcoming a), Coates (2012), Weatherson (ms), and Wedgwood (2012).

So, how do (ER), (ILC), and (PIF) generate a problem? Take any possible situation in which (i) and (ii) of (PIF) obtain. From (PIF)(i) and (ER), it follows that:

(A) S is rationally required to take D(p).

But from (PIF)(ii) and (ER) it follows that:

(B) S is rationally required to believe that her total evidence does not support taking D(p).

But this runs up against (ILC). For (A) and (B) rationally require a combination of attitudes that (ILC)(ii) prohibits. So when (PIF) obtains one is both *rationally required* and *rationally forbidden* to take an akratic combination of attitudes towards a proposition. Some refer to this puzzle as *the problem of misleading higher-order evidence*. I'll just call it *the PIF-Puzzle*.⁵

Existing solutions to this puzzle involve denying (PIF),⁶ denying (ILC),⁷ and endorsing the existence of dilemmas of rationality—i.e. circumstances where one cannot satisfy every demand of rationality.⁸ But these solutions are difficult to maintain for, as Worsnip (forthcoming) has persuasively argued, each of these solutions has costly drawbacks. Moreover, as Worsnip shows, there is at least one independently motivated and unexplored solution to this problem.

[[Another solution to this problems involves exploiting the distinction between propositional and doxastic justification. I've argued elsewhere that this offers an independently motivated and intuitive solution to the problem of misleading higher-order evidence. I also think it offers a solution to the

⁵For other formulations of this puzzles or puzzles very closely related to this see Littlejohn (forthcoming a) and Lasonen-Aarnio (forthcoming).

⁶Feldman (2005), White (2007, 120), Huemer (2011), Greco (2014), and Titelbaum (2015). Cp. Horowitz (2014).

⁷Coates (2012), Lasonen-Aarnio (forthcoming), Weatherson (ms), and Wedgwood (2012).

⁸This is how Worsnip reads Christensen (2007, 2010, 2013).

puzzle of misleading higher-order *apparent* evidence that I develop below. See my “How Doxastic Justification Helps Us Solve the Puzzle of Misleading Higher-Order Evidence” and the final comment of this paper for a bit more detail.]]

2 Separating Epistemic Reasons from Epistemic Rationality

Worsnip argues that when it comes to (ER) epistemologists have a lesson to learn from the literature on practical reason, where many philosophers distinguish *rationality* from *reasons-responsiveness*.⁹ For example, the kinds of cases that motivate this distinction are cases where *from the perspective of the agent* the balance of practical reasons support ϕ -ing, when in reality the balance of practical reasons support some activity that does not involve ϕ -ing. When an agent ϕ s in such a case, she acts in a practically rational way despite acting in a way that fails to be responsive to the balance of practical reasons. Worsnip sees (ER) as conflating these distinct normative statuses in the epistemic domain; in epistemology we need to differentiate *epistemic rationality* from *evidence-responsiveness*. Thus, while (ILC) is a genuine constraint on epistemic rationality, (ER) is not.

So Worsnip proposes we reject (ER) in favor of the following:

(ER*) If S’s total evidence supports D(p), then D(p) is the evidence-responsive attitude for S to take.¹⁰

Without (ER) the PIF-Puzzle dissolves. For without (ER) we cannot get (A) or (B) from (PIF)(i) and (ii). So the PIF-Puzzle is no puzzle at all.

⁹See Parfit (2001), Schroeder (2008, 2009, 2011), Gardner (2007), and Way (2009).

¹⁰Worsnip further cashes out evidence-responsiveness in the language of reasons: D(p) is an evidence-responsive attitude to take iff D(p) is supported by the balance of one’s *evidential (=epistemic) reasons*.

(ER*) together with (PIF) just shows us that thinkers can be in the unfortunate circumstance where evidence-responsiveness demands one *take* $D(p)$ while also *believing that one's total evidence doesn't support taking* $D(p)$. Following one's evidence in such cases will lead to irrationality, while obeying the dictates of rationality will lead to a failure to respond correctly to one's evidence. Such circumstances are troubling and undesirable, but not impossible. It's just part of the epistemic life of cognitively limited creatures who can be afflicted by misleading higher-order evidence.

It is a non-trivial virtue of Worsnip's solution that it finds independent motivation for rejecting (ER) in the literature on practical reason. For it makes the rejection of (ER) more than the best of a bad set of choices; rejecting (ER) is something we perhaps needed to do anyway.

3 Apparent Evidence and Coherence: the Conflict

Unfortunately the analogy with practical rationality comes back to bite. For even though practical rationality does not depend on what practical reasons there are, practical rationality does depend on what practical reasons there are *from the perspective of the agent*. Whiting (2014, 3) gives an example:

Tom is suffering from an illness that, if not treated immediately, will cause long-term damage. His doctor, Martha, holds a pill that she *believes* would kill Tom were he to take it. The many medical reports *she has read* and clinical trials she has conducted appear to confirm this. ... However, due to an unusual quirk in his physiology of which *no one is aware*, the pill will cure Tom.

According to the kind of view of rationality that has emerged in the literature on practical reasons, this sort of case is one where the balance of practical reasons support Martha giving the pill to Tom while it remains rational for

her not to give the pill. For what Martha has practical reason to do is grounded in, for example, the facts about the pill and Tom’s physiology; but the rationality of her action is grounded in her *perspective* on those facts—a perspective which is in fact mistaken. Just how one’s perspective on the facts grounds rationality is a matter of controversy. But what is not a matter of controversy is that it’s her perspective on the facts that determines what’s rational for her to do.¹¹

For convenience I’ll use the term ‘apparent reason’ in place of “reasons there are from an agent’s point of view.” And since Worsnip (forthcoming) analyzes evidence in terms of epistemic reasons, we can follow him and take *apparent epistemic reasons* to be *apparent evidence*.

Now if we’re to think about rationality as a general phenomenon that shares structural features across epistemic and practical (and certain other) normative domains as Worsnip suggests, then just as practical rationality tracks apparent practical reasons, epistemic rationality should track apparent epistemic reasons, i.e. apparent evidence. Thus we have:

APPARENT EVIDENCE REQUIREMENT (AER). If S’s total apparent evidence supports $D(p)$, then rationality requires of S that she takes $D(p)$.

There is no inconsistency in the rejection of (ER) and the acceptance of (AER) because (AER) concerns only apparent evidence, and apparent evidence isn’t necessarily actual evidence. Rejecting (ER) and endorsing (AER) represents a non-trivial shift in our epistemic ideology. For (AER) tethers rationality directly to apparent evidence, and only indirectly tethers rationality to actual evidence insofar as they overlap.

¹¹For example, the counterfactual analysis of practical rationality says (very roughly): it is practically rational for S to ϕ iff were what S believes true, it would give her a practical reason to ϕ . See Schroder (2007), Way (2009), Parfit (2001)). For a recent discussion of this analysis see Whiting (2014).

Other epistemologists have defended (AER) at length,¹² drawing both on the analogy with practical reasons and practical rationality as well as the way in which (AER) can help explain internalist intuitions about the positive epistemic standing of the beliefs of radically deceived thinkers (en-vatted brains and the like). And this seems right. For while the radically deceived cannot be responsive to actual evidence (so say many externalists), they can be responsive to apparent evidence, and it seems as if that later sort of responsiveness demands that one be credited with a rational belief. (AER) elegantly explains that intuition. But one needn't be an externalist to appreciate (AER). For so long as an internalist does not deny that we can sometimes be rationally misled about what our actual evidence supports, we will need (AER) or something like it to explain one's rational beliefs in such cases.¹³

At this point it's easy to see where things are heading. For if (PIF) is true, then we have good reason to think the following is true also:

POSSIBILITY OF APPARENT ITERATIVE FAILURE (APIF).

It is possible that:

- (i) S's total *apparent* evidence supports D(p), while
- (ii) S's total *apparent* evidence supports believing that her total actual evidence doesn't support D(p).

The reason for thinking (PIF) supports (APIF) is that the same kinds of paradigmatic examples that are used to support (PIF) also seem to support (APIF). Consider the following case where (PIF)(i) and (ii) obtain:

¹²Clayton Littlejohn (forthcoming b) and Kurt Sylvan MS. Cf. See Schroeder (2007: Ch. 1) and (2008). Note that each of these authors use different terms in drawing the distinction between reasons, apparent reasons, and rationality. For example, apparent reasons (Sylvan) are also named "excuses" (Littlejohn) and "subjective reasons" (Schroeder); while rationality is also named "excusability" (Littlejohn) and "substantive rationality" (Sylvan).

¹³For example, mentalism (Feldman and Conee 2004) seeks to ground epistemic reasons in one's mental states, but that's consistent with one lacking infallible access to what one's mental states are and what propositions one's mental states support.

Sherlock & Watson's Nephew. Sherlock is a detective who is famously good at assessing evidence. Sherlock has grudgingly granted Watson the favor of allowing his nephew to be his apprentice to see if he's got a future in detective work. Sherlock is repeatedly frustrated with Watson's nephew because he regularly makes mistakes in his assessment of what the evidence supports. The nephew's errors have been frequent enough to give the nephew good inductive evidence that he's bad at assessing the evidence. Sherlock has verbally confirmed this. While investigating their next crime scene Sherlock instructs the nephew to assess the evidence and come to a conclusion about what it supports. Among the pieces of evidence Watson's nephew comes to know of is a high definition video of the murder caught on a security camera. The recording provides conclusive (or near conclusive) evidence that the butler committed the murder. However, when Watson's nephew next encounters Sherlock, Sherlock tells the nephew that the evidence supports believing that the nanny is the murderer. However, Sherlock only draws this conclusion because he was never made aware of the video. And the nephew never discusses the video with Sherlock because he assumes Sherlock is already aware of it and he is too ashamed of his past failures to bring it up and ask Sherlock for further explanation.

Here is a case where the nephew's total actual evidence supports:

- (i) Believing that the butler did it, and (ii) believing that the total actual evidence does not support believing that the butler did it.

At least this is how it seems to advocates of (PIF).¹⁴ Advocates of (PIF) will, for example, hold that the nephew's knowledge of the video recording

¹⁴See Worsnip (forthcoming, section iv(b)) for a protracted defense of such cases in support of (PIF) as well as the other references in footnote 4.

is sufficient evidence to believe the butler is guilty. While it is Sherlock's testimony and the nephew's knowledge of his own bad track-record that is sufficient evidence to believe that the total actual evidence does not support believing that the butler did it. So the above example is one kind of example where (PIF)(i) and (ii) obtain.

It's a short step from such instances of (PIF) to similar instances of (APIF). For there can be overlap between apparent evidence and actual evidence, and if there is such overlap in PIF-cases then the puzzle we began with re-emerges. Of course, much turns here on what it takes to have sufficient apparent evidence.¹⁵ Here are three different kinds of conditions implied by prominent positions in the literature on apparent reasons:

(SAE: De Dicto) E is sufficient apparent evidence for S to take D(p) IFF it appears to S that E is sufficient actual evidence to take D(p).

(SAE: De Re) E is sufficient apparent reason for S to take D(p) IFF (i) E's truth would be sufficient actual evidence for S to take D(p), and (ii) it appears to S that E.

(SAE: Treating) E is sufficient apparent evidence for S to take D(p) IFF (i) it appears to S that E, (ii) S is attracted to treating E like a sufficient objective reason to take D(p), and (iii) this attraction manifests S's relevant reasons-sensitive competence, where a relevant reasons-sensitive competence is a competence to treat E-like considerations like objective reasons to take D(p)-like attitudes only if they are, when true, objective reasons to take D(p)-like attitudes.¹⁶

¹⁵I'm grateful to the Inquiry referee who prompted me to consider this in greater detail.

¹⁶These can be inferred from the following, more general claims discussed by Sylvan (2015):

(De Dicto) R is an apparent normative reason for S to ϕ iff it appears to S that R is an objective normative reason to ϕ .

Each view makes reference to ‘appearances’, which for present purposes may be understood to include both doxastic states (e.g. belief-like states) and non-doxastic states (e.g. seeming states).¹⁷

Both (SAE: De Re) and (SAE: Treating) imply that if there are PIF-cases like the case of Watson’s nephew then there are also corresponding APIF-cases. For Watson’s nephew’s actual evidence consists of, E1, the fact there is a video recording of the crime being committed by the butler and, E2, the fact that Sherlock has asserted that the total actual evidence supports believing that the nanny is the murderer. But this is not merely sufficient *actual* evidence for Watson’s nephew, it’s also sufficient *apparent* evidence according to (SAE: De Re) and (SAE: Treating). For if this is a PIF-case, then (SAE: De Re)’s condition (i) is fulfilled by the fact that E1 and E2 are sufficient actual evidence to believe the butler is the murderer but that the actual evidence doesn’t support believing that; and condition (ii) is fulfilled by the fact that Watson’s nephew knows E1 and E2.¹⁸

(SAE: Treating)’s conditions are likewise fulfilled: (i) is fulfilled by the fact that Watson’s nephew knows E1 and E2; (ii) is fulfilled by the fact that he’s is attracted to treating E1 like a sufficient objective reason to believe the

(De Re) R is an apparent reason for S to ϕ when (i) R’s truth would give S an objective normative reason to ϕ , and (ii) it appears to S that R.

(Treating) R is an apparent reason for S to ϕ iff (i) it appears to S that R, (ii) S is attracted to treating R like an objective reason to ϕ , and (iii) this attraction manifests S’s relevant reasons-sensitive competence, where a relevant reasons-sensitive competence is a competence to treat R-like considerations like objective reasons to do ϕ -like things only if they are, when true, objective reasons to do ϕ -like things.

Scanlon (1998: 25) and Kolodny (2005) appear to hold the de dicto view; Whiting (2014) endorses a modified version of it. Parfit (2001, 2011), Schroeder (2007), and Way (2009) appear to hold the de re view. Sylvan (2015) endorses a refined version of the treating view, however the refinements don’t impact our present discussion. So I’ll work with this simplified view.

¹⁷See Sylvan (2015) for further discussion of the relevance of including non-doxastic seeming states in one’s account of apparent reasons.

¹⁸I assume here that any adequate view of appearances will imply that: if S knows that P, then P appears to be the case to S.

butler did it and also attracted to treating E2 like a sufficient objective reason to believe that the total actual evidence doesn't support that belief; and, as is implicit in the case, (iii) is fulfilled by the fact that Watson's nephew's attractions are grounded in his relevant reasons-sensitive competences. Because of this, (SAE: Treating) also implies that if this is a PIF-case, it's also an APIF-cases.

Does (SAE: De Dicto) imply that if Watson's nephew is in a PIF-case then he's also in an APIF-case? Perhaps not. For part of what's at issue in thinking about apparent evidence for S is S's first-person point of view, and some think that the first-person point of view is not a point of view from which we can appreciate PIF-cases. For example, Worsnip (forthcoming, section 3) writes:

...the state of believing that one's evidence does not support believing p, but nevertheless believing p, is harder to make sense of, at least from a first-personal perspective. It amounts to saying 'I have nothing that gives any adequate indication to me that p is the case; nevertheless, p is the case'. *First-personally, these states do not seem capable of withstanding serious reflection.* And third-personally, while we can imagine such agents, in describing and explaining them *we reach for some story involving self-deception or a failure to recognize their own mental states.*¹⁹

If this is right, then (SAE: De Dicto) is unlike (SAE: De Re) and (SAE: Treating) in that it offers a way of endorsing (PIF) while rejecting (APIF). I'll return to this issue.

So here is where we are at. The problem from the previous section involving (ER), (PIF), and (ILC) seemed to be neatly resolved by Worsnip.

¹⁹See also Gibbons (2014) and Foley (2001). Worsnip's reference to the "first-person perspective" is to be understood in terms of one's *beliefs* about one's own situation, and not the general sense we began with when we introduced the general notion of apparent evidence.

But the puzzle just resurfaces with (AER), (ILC), and certain prominent views of apparent reasons that, together with (PIF), imply (APIF). For from (APIF)(i) and (AER) it follows that:

(A) S is rationally required to take D(p).

But from (APIF)(ii) and (AER) it follows that:

(B) S is rationally required to believe that her total actual evidence does not support taking D(p).

And as we've already seen, (A) and (B) conflict with (ILC)(ii). So the very same puzzle of rationality remains even if Worsnip is correct that we must reject (ER) and distinguish *rationality* from *being evidence-responsive*.

4 Moving Forward

One way of responding to the revived PIF-Puzzle is to follow Worsnip's divide-and-conquer strategy: just distinguish *rationality* from being *apparent-evidence-responsive* and thus reject (AER) and introduce something like:

(AER*) If S's total apparent evidence supports D(p), then D(p) is the apparent-evidence-responsive attitude for S to take.

This is what some theorists are getting at when they draw a distinction between *substantive rationality* and *structural rationality*—where being substantively rational has to do with being apparent-evidence-responsive and being structurally rational has to do with avoiding attitudinal conflicts irrespective of whether or not it would be substantively rational for one to take any of the conflicting attitudes (cf. Broome 2013; Scanlon 1998; Sylvan MS). This kind of distinction might block the revived PIF-Puzzle.²⁰

²⁰I'm grateful to the Inquiry referee who prompted me to consider this issue.

Accordingly, it would seem like there are two solutions to the revived PIF-Puzzle. First one could draw the distinction between substantive and structural rationality and endorse something like (AER*), while rejecting both (AER) and (ER) when the term ‘rationality’ is taken to refer to structural rationality. Second, as discussed above, one could adopt (SAE: De Dicto) and maintain that it cannot appear to oneself that one is in a PIF-case “in a way that withstands serious reflection,” as Worsnip suggests. Both responses seem to block APIF-cases, and without such cases there is no revived puzzle.

But there are reasons to worry about both ways of responding to the revived puzzle. For drawing the distinction between substantive and structural rationality has an ad hoc flavor about it once we concede that there are extremely clever, honest, and well-informed people who self-consciously maintain beliefs that violate paradigmatic principles of structural rationality. For example, consider the following claims:

(Contradiction) P is true, and P is false.

(Dilemma) I am rationally required to believe P and to not believe P.

(Misleading) P, and my total actual evidence insufficiently supports believing P.

Pre-theoretically, the idea that one could rationally believe (Contradiction) or (Dilemma) is at least as counter-intuitive as the idea that one could rationally believe (Misleading). This pre-theoretic judgement has led some to think that structural rationality is, at least in part, a matter of avoiding such combinations of attitudes. But it’s hard to spend time in conversation with, say, J. C. Beall or Graham Priest, read their books and associated articles, and then draw any other conclusion than that they are extremely clever, honest, and well-informed people who *rationally* and self-consciously believe (Contradiction) and (Dilemma) on the basis of their, at least, *apparent* evi-

dence.²¹

But if Beall and Priest can rationally believe (Contradiction) and (Dilemma) on the basis of their apparent evidence, then it's hard to tell what relevant notion of rationality the stipulative term 'structural rationality' refers to if even these masters of the structural relations that obtain between propositions fail to be structurally rational. So rejecting (AER) and embracing a division between substantive and structural rationality threatens to be more of an illusive shift in topic than a genuine solution to the PIF-Puzzle. At the very least, those who think that there is an important notion of structural rationality that can be leveraged to reject (AER) and solve the PIF-Puzzle need to justify their position vis-a-vis the present concern.²²

These worries about structural rationality likewise affect those who might seek a way out of the PIF-Puzzle by taking (SAE: De Dicto) and maintaining that one cannot view oneself as being in a PIF-case "in a way that withstands serious reflection." For Beall and Priest are paradigms of people who rationally *and self-consciously* believe (Contradiction) and (Dilemma) in a way that withstands serious reflection; and in describing them as such we do not (or at least should not) reach for "some story involving self-deception or a failure to recognize their own mental states." But if it's possible for them to rationally believe (Contradiction) and (Dilemma) on the basis of their apparent evidence, then it seems like the same should be possible for (Misleading). So even if (SAE: De Dicto) is correct, it's not clear that it offers a way out of the PIF-Puzzle.

In the introduction I began by citing the plausibility of (ER) and I later acknowledged the plausibility of (AER). Both cannot be correct—assuming that apparent evidence and actual evidence are not coextensive. So which is right? I leave it to the reader to decide. For it does not matter for the purposes of resolving the PIF-Puzzle. What matters is that *either (ER) is*

²¹See Beall (2009) and Priest (2005).

²²For further development of the present concern with 'structural rationality' see Lasonen-Aarnio (MS).

true or (AER) is true. One can reject either one of these without disfiguring our notion of rationality beyond recognition. But rejecting both is, arguably, more than our ordinary way of thinking about rationality can bear. And being stuck with one, whichever it is, is enough to enable the PIF-Puzzle.

So where do we go from here? In the end, there may be no ideal solution to the PIF-Puzzle. Recall the existing responses to this puzzle: rejecting (PIF), rejecting (ILC), or endorsing the existence of dilemmas of rationality. Going a step further than Worsnip did, we can add two more possible responses. First, the rejection of both (*ER*) and (*AER*) together with the introduction of some theory of ‘structural rationality’ that is capable of plausibly responding to the challenge above. Second, endorse some view of apparent evidence that doesn’t entail the following: if there are PIF-cases, then there are APIF-cases. The thing to observe is that each of these solutions has non-trivial costs. In the end, there may be no good solution to the PIF-Puzzle.

[[One attempt to solve the PIF-puzzle that seems to avoid all the counterintuitive consequences of the previously noted attempts involves drawing a distinction between propositional and doxastic *justification* (construed in terms of reasons-responsiveness). If that works, and if we can extend the target notions of propositional and doxastic justification to *rationality* (construed in terms of *apparent* reasons-responsiveness), then there is a promise of an unproblematic solution to the APIF-puzzle discussed in this paper. See Silva’s (2017) “How Doxastic Justification Helps Us Solve the Puzzle of Misleading Higher-Order Evidence” for how the difference between propositional and doxastic justification can help here.]]

References

- Beall, J.C. 2009. *Spandrels of Truth*. Oxford: Oxford University Press.
- Broome, J. 2013. *Rationality Through Reasoning*. Chichester:

Wiley-Blackwell.

Christensen, D. 2007. "Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals," *Oxford Studies in Epistemology*, 2: 3-31.

Christensen, D. 2010. "Higher-Order Evidence," *Philosophy and Phenomenological Research*, 81/1: 185-215.

Christensen, D. 2013. "Epistemic Modesty Defended," in Christensen & Lackey (2013).

Coates, A. 2012. "Rational Epistemic Akrasia," *American Philosophical Quarterly*, 49/2: 113-124.

Conee, E. and R. Feldman. 2004. *Evidentialism*. Oxford University Press.

Elga, A. 2005. "On Overrating Oneself...and Knowing It," *Philosophical Studies* 133: 115-124.

Feldman, R. 2005. "Respecting the evidence," *Philosophical Perspectives*, 19: 95-119. Blackwell, Oxford: *Epistemology*.

Gardner, J. 2007. *Offenses and Defenses*. Oxford University Press.

Greco, D. 2014. "A puzzle about epistemic akrasia," *Philosophical Studies* 167/2: 201-219.

Horowitz, S. 2014. "Epistemic Akrasia," *Noûs* 48/4: 718-744.

Huemer, M. 2011. "The Puzzle of Metacoherence," *Philosophy and Phenomenological Research* 82/1: 1- 21.

Priest, G. 2005. *Doubt Truth to be a Liar*. Oxford: Oxford University Press.

Lasonen-Aarnio, M. 2014. "Higher-Order Evidence and the Limits of Defeat," *Philosophy and Phenomenological Research* 88:

314-345.

Lasonen-Aarnio, M. Forthcoming. "Enkrasia or evidentialism? Learning to love mismatch," *Philosophical Studies*.

Lasonen-Aarnio, M. MS. "Coherence." University of Michigan.

Littlejohn, C. Forthcoming a. "Stop Making Sense? A Puzzle About Epistemic Rationality," *Philosophy and Phenomenological Research*.

Littlejohn, C. Forthcoming b. "A Plea for Epistemic Excuses," in F. Dorsch and J. Dutant (ed.), *The New Evil Demon*. Oxford University Press.

Parfit, D. 2001. "Rationality and Reasons" in *Exploring Practical Philosophy*. Aldershot: Ashgate.

Scanlon, T. M. 1998. *What We Owe to Each Other*. Cambridge MA: Harvard University Press.

Schroeder, M. 2007. *Slaves of the Passions*. Oxford: OUP.

Schroeder, M. 2008. "Having Reasons." *Philosophical Studies* 139: 57–71.

Schroeder, M. 2009. "Means-End Coherence, Stringency, and Subjective Reasons." *Philosophical Studies* 143: 223–248.

Schroeder, M. 2011. "What Is It To 'Have' A Reason?" in *Reasons for Belief*. Cambridge: CUP.

Sylvan, K. 2015. "What apparent reasons appear to be," *Philosophical Studies* 172/3: 587-606.

Sylvan, K. MS. "On Divorcing the Rational and the Justified in Epistemology," Dissertation.

Titelbaum, M. 2015. "Rationality's Fixed Point (or: In Defense of Right Reason)," *Oxford Studies in Epistemology*, 5.

- Way, J. 2009. "Two Accounts of the Normativity of Rationality." *Journal of Ethics and Social Philosophy* 4/2: 1–8.
- Weatherson, B. MS. "Do Judgments Screen Evidence?," draft manuscript, University of Michigan.
- Wedgwood, R. 2012. "Justified inference," *Synthese* 189/2: 273–295.
- White, R. 2007. "Epistemic Subjectivism," *Episteme* 4/1: 115–129.
- Whiting, D. 2014. "Keeping Things in Perspective: Reasons, Rationality, and the A Priori," *Journal of Ethics and Social Philosophy* 8/1: 1–23.
- Worsnip, A. Forthcoming. "The Conflict of Evidence and Coherence," *Philosophy and Phenomenological Research*.