AGAINST COHERENCE

Darren BRADLEY

ABSTRACT: Some philosophers argue that coherence is a normatively significant type of rationality, over and above substantive rationality. The most detailed and substantive arguments are given by Alex Worsnip (2021, 2022). In this paper I will criticize his arguments. We are left with the thesis that the only type of rationality we need is substantive rationality.

KEYWORDS: structural rationality, substantive rationality, coherence, deliberation

1. Introduction

Many philosophers distinguish *substantive* rationality from *structural* rationality (also called coherence). When an agents' beliefs fit the evidence they are substantively rational; when an agent's beliefs fit together in the right way they are structurally rational / coherent. Some philosophers argue that coherence is a normatively significant type of rationality, over and above substantive rationality. I disagree. I will argue that the only type of rationality we need is substantive rationality.

Let's distinguish three positions:

Substantive reasons monism
The only constraints on rationality are substantive reasons (Kolodny 2005, Kiesewetter 2017 and Lord 2018)

Coherence constraint monism
The only constraints on rationality are coherence constraints (Broome 2013)

Dualism
Rationality involves both substantive reasons and coherence constraints (Worsnip 2021)

I defend substantive reasons monism. I will focus on the recent book-length treatment of Worsnip (2021) which offers the most detailed and sophisticated defence of dualism, and also draw on his (2022) paper. Worsnip and I agree that there are substantive reasons, so I will take these as given. I can find two arguments in the book for coherence constraints on rationality, and will argue that neither succeed.

Worsnip's book is part of the literature on reasons in meta-ethics and epistemology, but the question of whether there are coherence constraints on

rationality has even broader implications. A foundational assumption of Bayesianism is that all rational agents are certain of tautologies (e.g. p or not p) (Lin 2023 section 1.9). This amounts to a coherence constraint on rationality. Relatedly, one current of logical exceptionalism (Williamson 2007, Hjortland 2017, Read 2019) says that all rational agents are certain of all logical truths, which suggests that the logical truths are coherence constraints on rationality. By contrast, Quine (1951) famously argued that all statements were revisable, suggesting a rejection of any coherence constraints on rationality. The arguments I give against coherence support Quine's side of these debates.

In the next section I elaborate on the distinction between coherence and substantive rationality. In section 3 I argue that coherence constraints are not needed to structure deliberation. In section 4 I offer some cases intended to show that coherence is not a normatively significant type of rationality. These are analogous to the cases Worsnip offers, and which I discuss in section 5. Section 6 concludes.

2. Substantive Rationality vs Coherence

Let's start by getting clear on the distinction between substantive rationality and coherence. The best way to get a grip is with examples. Many of our beliefs are rational or irrational due to our circumstances. For example, if an agent's total evidence is that there seems to be a red object in front of them, then it is *substantively* rational for them to believe that there is a red object in front of them. By contrast, there are some beliefs such that their rationality or irrationality is not due to our circumstances. Suppose the agent believes there is a red object in front of them and also not a red object in front of them. There seems to be something wrong with this agent that is independent of their circumstances. Whether they seem to see a red object or not, they should not believe that there is both a red object and not a red object in front of them. In doing so, they have *incoherent* attitudes.[1] They are irrational in a specific sense that Worsnip calls *structural* irrationality.

Moving from examples to a more general understanding, Worsnip offers three hallmarks of (in)coherence:

> 1. Making a judgment of the [(in)coherence] of a set of attitudes does not require information about the circumstances or evidence of the agent in question, whereas

[1] I leave open the question of exactly which attitudes are incoherent. Worsnip (2021, 133) says agents have a disposition to revise incoherent attitudes while Lee (2024) says agents with incoherent attitudes have incompatible commitments. Daoust (2023) argues that there is no clear meaning for 'coherence'.

> making a judgment of the substantive rationality of an attitude typically does.
>
> 2. Judgments about the [(in)coherence] of attitudes can be made in abstraction from disagreements about what is worth doing, which considerations constitute reasons for which responses, and so on—whereas judgments about substantive rationality cannot.
>
> 3. Indeed, in saying which combinations of attitudes are [incoherent], we can point to patterns of [coherence], using variables and schematic symbols, and without even fully specifying the content of the attitudes in question. (7-8)

Having got clear on the distinction, why should we posit coherence constraints on rationality? Worsnip argues that coherence constraints are needed to capture two verdicts:

> i) 'that incoherence is a further kind of irrationality, over and above … substantive irrationality' (59) and
>
> ii) that coherent agents are 'in one sense more rational than their incoherent counterparts' (60)

I will argue against both (i) and (ii). In slogan form: substantive rationality is the only rationality.

One complication is that Worsnip argues for dualism about both practical rationality (relevant to reasoning that results in actions) and theoretical rationality (relevant to reasoning that results in beliefs). In this paper I focus purely on theoretical rationality.[2] My claim is that there is no need for coherence constraints in theoretical rationality. A fuller slogan: substantive rationality is the only theoretical rationality.

Terminology in this area is notoriously tricky. To get clear on one possible source of confusion, consider the following helpful comment:

> [C]oherence, or structural rationality, is a … normatively significant phenomenon ... When I say that it is genuine, I mean that being coherent genuinely amounts to a kind of rationality. (ix)

Worsnip is arguing that structural rationality is a kind of rationality. This looks to be true by definition, but Worsnip does not intend it to be true by definition. That's why I prefer to talk about *coherence* rather than structural rationality. We can ask whether coherence constraints are a kind of rationality, or normatively significant, or provide genuine reasons.[3] I take these locutions to be equivalent. Compare: Both sides agree that substantive reasons are a kind of rationality, or

[2] I'm inclined to think we do need structural constraints for practical rationality (see Scanlon 2003).
[3] I secretly find this talk of 'normative significance' and 'genuine reasons' highly problematic, but I will put aside such qualms for this paper.

normatively significant, or provide genuine reasons; both sides agree that etiquette constraints are *not* are a kind of rationality, nor normatively significant, nor provide genuine reasons. Worsnip argues that coherence constraints are like substantive reasons; I argue that they are more like etiquette constraints.

I follow Worsnip's terminology in taking 'rationality' to be normatively significant / provide genuine reasons. Unfortunately it is also natural to say that it is rational to make some logically correct inference, but we do not want to pre-judge whether doing so is normatively significant / fits genuine reasons. So where there is a risk of confusion I will talk about who is *doing better*, or whether an inference is *good,* in order to reflect normative significance / genuine reasons.

3. Deliberation

I will argue that we have no need to posit coherence constraints on rationality. Most of the contemporary literature on the topic is responding directly or indirectly to Kolodny (2005), who rejects coherence constraints. One of his arguments is that coherence constraints cannot play a role in deliberation, so cannot be reasons (given that reasons must be things we can reason *with*). Worsnip replies with a detailed account of how coherence constraints can play a role in deliberation, aiming to show that coherence constraints are reasons.[4] I will argue that coherence constraints do not play the role in structuring deliberation that Worsnip wants.

Part of Worsnip's reply to Kolodny suggests a reason for positing coherence constraints. Specifically, Worsnip argues that coherence constraints *structure* deliberation as opposed to entering into deliberation. Consider an agent deliberating about whether to believe p, not p or neither. Worsnip writes:

> … when we deliberate, we often don't deliberate about individual attitudes (or their objects) in isolation. (252)

> Specifically, letting *d* be the proposition that a Democrat will be US President in 2025, I treat both {believing that the evidence supports believing *d*, not believing *d*} and {believing that the evidence doesn't support believing *d*, believing *d*} as off-limits. This means that I am effectively deliberating between only two options: {believing that the evidence supports believing *d*, believing *d*} and {not believing that the evidence supports believing *d*, not believing *d*}. So I make up my mind about both questions together simultaneously and in concert. (255)

[4] The dialectic is in fact far more complicated than this. Worsnip offers three readings of Kolodny's argument, but I think none of them bring into focus Kolodny's main point that we cannot reason using coherence considerations. I think we can reason using coherence considerations (which are derived from substantive rationality) but I can't defend any of this here.

Worsnip holds that substantive rationality cannot explain this, only coherence constraints can. Notice that Worsnip's argument assumes that we should always treat {believing that the evidence supports p, not believing p} as off-limits. Let's assume for the sake of argument that we should.

Do we need coherence considerations to explain why we should always treat {believing that the evidence supports p, not believing p} as off-limits? Not usually, for in most cases either:

> i) the evidence supports believing p, in which case it is substantively rational to believe that the evidence supports p and to believe p, or
>
> ii) the evidence supports believing not p, in which case it is substantively rational to believe that the evidence supports believing not p and to believe not p, or
>
> iii) the evidence supports believing neither p nor not believing p, in which case it is substantively rational to believe the evidence supports neither believing p nor not believing p, and to neither believe p nor not p.[5]

None of these options allow {believing that the evidence supports p, not believing p}. So rather than appealing to coherence considerations, we can appeal to background knowledge that i, ii or iii hold to explain why we should treat {believing that the evidence supports p, not believing p} as off-limits.

Worsnip offers two compatible responses. In his (2022) paper, he argues that even if the disjunction of i, ii or iii is true in all cases, coherence considerations are not redundant; in his (2021) book, he denies that the disjunction of i, ii or iii are true in all cases. Let's take these in turn.

In his (2022) paper, he writes:

> Let's suppose, for the sake of argument, that the claim [that] your substantive reasons never permit an incoherent set of attitudes is correct. If this is so, then it's true that coherence requirements don't impose any additional demands on your attitudes, over and above those imposed by your substantive reasons. Nevertheless, the present account isolates a further role that coherence considerations should play in your deliberation. Suppose you took incoherent combinations of attitudes seriously in your deliberation, and did not focus deliberation solely on the coherent combinations. Even if you were to then go on to rule out each of the incoherent combinations on the grounds that they contain attitudes for which you lack sufficient substantive reasons, and thus ultimately end up with coherent attitudes, there would be something defective about your deliberation *qua* deliberation due to your having taken the incoherent combinations seriously to start with. (408)

[5] I've assumed Uniqueness for simplicity. A permissivist analysis will be more complicated but will still not allow {believing that the evidence supports p, not believing p}.

Worsnip is arguing that we should posit coherence constraints in order to rule as defective deliberation that takes seriously incoherent combinations. However, it is far from obvious that there is anything defective about deliberation that takes seriously incoherent combinations. Let me fill this out a little by thinking about the role deliberation plays for ideal and non-ideal agents. Let an ideal agent be an agent that always instantly arrives at the correct belief (or credence, or decision). I have in mind Bayesian agents who satisfy the axioms of probability and instantly conditionalize on any new evidence. They are more like angels or Gods than humans. What role does deliberation have in their thought? Plausibly, none. Deliberation is a process of thinking things through and arriving at a verdict. The closest thing to deliberation that these ideal agents go through is instantaneously forming the correct belief/credence/decision. But that is not what we mean when we talk about deliberation. Deliberation is very much an activity of non-ideal agents like humans who need to take their time thinking things through, aware that they might well be mistaken. Deliberation can involve mistakes, go down blind alleys and arrive at unjustified beliefs/credences/decision—but that doesn't make it defective. Similarly, I don't think we should say that deliberation that takes seriously incoherent combinations of attitudes is defective. And if deliberation that takes seriously incoherent combinations of attitudes is not defective then there is no need to posit coherence constraints to rule it out.

Let's move on to Worsnip's second response. In his (2021 section 3.5.7) book Worsnip defends the possibility that none of i, ii or iii hold, and that instead:

> iv) the evidence supports believing that the evidence supports p and the evidence supports not believing p

This would be a case of rational *akrasia*. In such a case substantive rationality allows {believing that the evidence supports p, not believing p}, and so, according to Worsnip, we need coherence considerations to explain why this combination is off-limits.

But we have now arrived at a very odd set of views—that one should always treat {believing that the evidence supports p, not believing p} as off-limits yet it is possible for the evidence to support believing that the evidence supports p and for the evidence to support not believing p. I would suggest that, if you really think that it is possible for the evidence to support believing that the evidence supports p and for the evidence to support not believing p, then you should *not* treat {believing that the evidence supports p, not believing p} as off-limits.

Worsnip concedes that he is committed to an odd set of views and offers some ways to cushion the blow (section 8.8.4). I'm not going to take a stand on the degree to which he succeeds. My question is: why bother? Why maintain that one should

always treat {believing that the evidence supports p, not believing p} as off-limits? Isn't it possible that there is some situation where that is the best thing to believe? It seems to me at least an open question whether to treat {believing that the evidence supports p, not believing p} as off-limits, especially given that numerous philosophers defend the rationality of akrasia.[6] So in the absence of a strong argument that we should treat {believing that the evidence supports p, not believing p} as off-limits, we are left without a strong argument that coherence constraints are needed to structure deliberation.

4. Irrational Coherence

In this section I'll consider some cases which are analogous to those Worsnip discusses and argue that we find no support for coherence constraints.

Suppose two agents have the same irrational belief, p. The only difference between them is that the first agent has made many inferences from p and the second has made no inferences from p. Worsnip claims that the first agent is more rational in at least one respect. But I feel no pull to say that the first agent is more rational in any respect.

Indeed it is very plausible that the agent who makes many inferences from p is less rational simpliciter, for they have let the irrationality of p percolate out to the rest of their beliefs. Given the plausible assumption that any belief inferred from an irrational belief is also an irrational belief, the agent who makes many inferences from p has many more irrational beliefs than the agent who doesn't. It is odd to conclude that these irrational beliefs make the agent more rational in any respect than their counterpart who does not have such irrational beliefs. Quite the reverse, it seems like the extra irrational beliefs must make the agent less rational.

Let's go through a couple of examples.

Preacher and Post-doc

It is 1989, and Oliver is a talented post-doc in a respected biology department. He rationally believes that life on Earth evolved by natural selection. However, he has recently come under the influence of a charismatic preacher who says that the world was created by God in 4004 BC. Oliver irrationally believes that the preacher says true things, but he is unable to give up his belief that life on Earth evolved by natural selection. He just lives with the incoherence.

By 1990, Oliver's beliefs are fully coherent with the belief that the world was created by God in 4004 BC. As a result, he irrationally disbelieves that life on Earth

[6] See Arpaly (2003), Horowitz (2013), Christensen (2016, 2021, 2022,) Pryor (2018), Hawthorne, Isaacs & Lasonen-Aarnio (2021).

evolved by natural selection.

This looks like a terrible change. Oliver in 1990 is less rational in all ways than in 1989. The rational belief has been jettisoned because it conflicted with an irrational belief. This irrational belief has done a lot more damage by percolating out into the rest of Oliver's beliefs and causing him to give up rational beliefs.

But Worsnip cannot give these verdicts. According to Worsnip (2021, 50), coherent agents are 'in one sense more rational than their incoherent counterparts'. Oliver in 1990 is coherent, and so is more rational in one sense than he was in 1989.

Let me emphasize two points here. First, even if 1990-Oliver is more *coherent* than 1989-Oliver, the disagreement is whether this amounts to a type of *rationality*. How are we to decide whether coherence is a type of rationality? Worsnip's method is to appeal to cases and intuitions. We'll look at the case Worsnip offers below. For now, Worsnip would need to say that we intuitively judge that 1990-Oliver is more rational, in at least one sense, than 1989-Oliver. But I do not share this intuition. It seems to me that 1990-Oliver is simply doing worse than 1989-Oliver when it comes to rationality.

But perhaps we can dig a bit deeper into these intuitions. This leads to the second point. Recall Worsnip will claim that there is some specific type of rationality on which 1990-Oliver scores higher than 1989-Oliver i.e. coherence. Worsnip might concede that 1990-Oliver loses points for each belief that conflicts with his evidence, but gains points for each belief that coheres with 'God created the world in 4004 BC'.

But I see no point in positing this dimension of rationality. For the points 1990-Oliver gains for having coherent beliefs will be cancelled out by losing points for having beliefs which conflict with the evidence. Let's see why. Suppose 1990-Oliver does have greater coherence than 1989-Oliver in virtue of the beliefs he has inferred from his belief that God created the world in 4004 BC. Consider one of these beliefs – that dinosaur fossils were never parts of living creatures. Notice that this belief conflicts with 1990-Oliver's evidence, which strongly confirms that they were. 1990-Oliver is *substantively irrational* for having this belief which conflicts with his evidence. And it seems to me that the cost of this substantive irrationality will balance out the benefit of any coherence we might posit. So I see no point positing coherence constraints on rationality.

I can think of two ways Worsnip might respond. First, he might deny that the points gained from the coherent beliefs are always cancelled out by the points lost for having beliefs that conflict with the evidence. Second, he might argue that we need to posit points *gained* for coherent beliefs to cancel out the points *lost* for having beliefs which conflict with the evidence.

Starting with the first, how might one deny that the points gained from the coherent beliefs are always cancelled out by the points lost for having beliefs that conflict with the evidence? Presumably the agent might gain *more* points for coherent beliefs than are lost for beliefs that conflict with the evidence. The net result would be that the agent would become overall more rational by making inferences from beliefs that conflict with the evidence. But the example of 1989-Oliver and 1990-Oliver was designed to show that this is not the case. It does not seem like 1990-Oliver becomes overall more rational by inferring further beliefs from his belief that many world leaders are reptilian shape-shifters.

On the second response, Worsnip might argue that we need to posit points *gained* for coherent beliefs in order to cancel out the points *lost* for having beliefs which conflict with the evidence. The thought might be that when 1990-Oliver makes an inference from his belief that God created the world in 4004 BC he doesn't *only* lose points for having beliefs which conflict with the evidence. For there is something good about making correct (i.e. deductively valid and inductively good) inferences, and we need coherence constraints to allow for the goodness of making correct inferences.

But I don't see that there is anything good about making correct inferences. I agree that there is something good about making correct inferences *from rational beliefs*; but I don't see anything good about making correct inferences from irrational beliefs.

Let's consider one final way an objector might disagree. They might argue that there *is* something rational about making correct inferences as follows. We need to compare an agent who makes correct inferences from their beliefs to an agent who makes incorrect inferences from their beliefs. Let's introduce a third character, Charlie, who also believes that God created the world in 4004 BC despite strong evidence to the contrary. Worse, Charlie is bad at making inferences (perhaps he engages in confirmation bias, availability bias and wishful thinking), so he infers further beliefs that do not fit with his evidence, nor do they follow from his prior beliefs. To fill this out, suppose that both 1990-Oliver and Charlie have identical evidence. Both believe that God created the world in 4004 BC and a host of beliefs that follow from it. But Charlie also has a host of beliefs that *don't* follow from it! So, for example, suppose Charlie believes that God's favourite colour is aqua-marine (based on some fallacious inference from a biblical passage). We can all agree that Charlie is even worse, rationally speaking, than 1990-Oliver. But why? According to the objector, because 1990-Oliver has made *correct* inferences from his (irrational) beliefs and Charlie has made *incorrect* inferences from his (irrational)

beliefs. So there must be something good about making correct inferences, even from irrational beliefs.

In response, I claim that Charlie is less *substantively* rational that Oliver-1990. We can see why Charlie is doing so badly simply by comparing his beliefs with his evidence. For Charlie must be taking[7] his evidence to support that God's favourite colour is aqua-marine, while 1990-Oliver does not. 1990-Oliver's only mistake is to take his evidence to support that that God created the world in 4004 BC. If it did, 1990-Oliver would be right to make the many inferences he has made. By contrast, Charlie has made the additional mistake of taking his evidence to support that God's favourite colour is aqua-marine. Charlie is mistaken concerning what the evidence supports. This explains why Charlie is worse off than 1990-Oliver—Charlie is less substantively rational than 1990-Oliver.

Before we get to Worsnip's example in the next section, it might be worth modifying our example to make it more similar to Worsnip's. Our example involves one agent at two times, while Worsnip's involves two agents at one time. So we can adapt our example:

> **Preacher and Two Post-docs**
>
> Ashley and Glen are talented post-docs in a respected biology department.
>
> Ashley rationally believes that life on Earth evolved by natural selection, but has come under the influence of a charismatic preacher who says that the world was created by God in 4004 BC. Ashley irrationally believes that the preacher says true things, but he is unable to give up his belief that life on Earth evolved by natural selection. He just lives with the incoherence.
>
> Glen has fully bought into the preacher's teaching, and irrationally believes that the world was created by God in 4004 BC and everything that follows from it e.g he irrationally disbelieves that life on Earth evolved by natural selection.

I think our judgment will be similar to the temporal case. Glen seems to be rationally worse off in all ways than Ashley. Not only does he irrationally believe that the world was created by God in 4004 BC, he now irrationally believes things that follow from it.

[7] This sense of taking is based on Boghossian's (2014) use of what he calls the *taking condition* on inference. Boghossian is talking about inferences from a belief to another belief, but something similar is needed for the move from evidence to belief.

5. Worsnip's Example

Let's apply the points made in the previous section to Worsnip's example. Worsnip introduces two characters as follows[8]:

> Consider, for example, Disbelieving Derrick, who believes that he is Superman, and believes that Superman can fly. In believing that he's Superman, Disbelieving Derrick has a belief that flies in the face of his evidence, his evidential reasons for belief (or so we may safely assume). Moreover, it's very natural to describe this belief as irrational. But now let us add another piece of information about Disbelieving Derrick: he believes he can't fly … Intuitively, when we add this further piece of information, we can now see that Disbelieving Derrick is irrational in a second respect: specifically, he has inconsistent beliefs …
>
> To sharpen the intuitive point still further, we can compare Disbelieving Derrick to his brother Believing Billy. Like Disbelieving Derrick, Believing Billy believes that he is Superman and that Superman can fly, but unlike Disbelieving Derrick, Believing Billy believes that he can fly. There's a clear sense in which Believing Billy is even less rational than Disbelieving Derrick, since he has two beliefs that go dramatically against his evidence (viz. that he is Superman, and that he can fly), where Disbelieving Derrick has only one (viz. that he is Superman). But there is also a clear sense in which Believing Billy is more rational than Disbelieving Derrick, since his beliefs cohere in a way that Disbelieving Derrick's don't: there's no inconsistency in his beliefs (and, indeed, his belief that he can fly is the logical consequence of his other two beliefs). Again, we can recognize both of these senses by saying that Disbelieving Derrick is less substantively irrational than Believing Billy, but Believing Billy is less [incoherent] than Disbelieving Derrick. (5-6)

This case is analogous to *Preacher and Two Post-docs*, and I think the same verdicts apply. In my view, there is no sense in which Believing Billy is less irrational than Disbelieving Derrick; Disbelieving Derrick is more rational *simpliciter*.

Believing Billy believes he can fly as he infers it from his belief that he is Superman (and that Superman can fly). But Billy's belief that he can fly conflicts with his evidence that he cannot fly. The more inferences he makes from his irrational belief that he is Superman, the more irrational he becomes. Equivalently, the more beliefs Billy forms which are in conflict with his evidence, the more irrational he becomes.

Analogous points hold of Disbelieving Derrick. We are told that Disbelieving Derrick believes he can't fly . Presumably Disbelieving Derrick believes he can't fly because his evidence indicates that he can't fly. As such, the belief that he cannot fly increases his rationality. It conflicts with his irrational belief that he is Superman,

[8] Worsnip calls them Tom and Tim. I re-name them to make them easier to identify.

but further beliefs which fit with his belief that he is Superman would make him less rational, not more.

I concede that intuitions are less clear in this case, perhaps because it is so difficult to imagine a person who believes that they are Superman. But to the extent that our intuitions are unclear in this case, we should default to our intuitions about the cases from section 4.

6. Conclusion

I have considered two arguments in Worsnip's book for the claim that there are coherence constraints on rationality. The first argument is that coherence constraints are needed to structure deliberation. I argued that we have little reason to think that deliberation should be so structured. The second argument is based on cases. I argued that the cases do not make a strong argument for positing coherence constraints.

Let me finish with a point about the dialectic. I do not claim to have shown that there are no coherence constraints. My aim has been to respond to Worsnip's arguments that there are. Worsnip is positing something new,[9] and if you posit something new the burden is on you to explain why we need it. Furthermore, positing both substantive and structural rationality raises tricky questions about how the two interact. Are they comparable? Can one be traded off against the other? In the absence of strong arguments for coherence constraints, I suggest that we avoid these questions by endorsing the simpler and more elegant theory that there are only substantive reasons to believe.[10]

References

Arpaly, N. 2003. *Unprincipled Virtue: An Inquiry into Moral Agency*. Oxford University Press, USA.

Bradley, D. 2018. "Philosophers Should Prefer Simpler Theories." *Philosophical Studies* 175(12): 3049-3067

Boghossian, P. 2014. "What Is Inference?" *Philosophical Studies* 169(1): 1-18. https://doi.org/10.1007/s11098-012-9903-x

Broome, J. 1999. "Normative requirements." *Ratio* 12(4): 398-419.

[9] I don't mean new historically. Others have defended coherence constraints. I mean new in the logical/dialectical sense. Worsnip's opponent is denying that coherence constraints exist; relative to this opponent, Worsnip is adding something new.

[10] I'm grateful to the John Templeton Foundation (grant 62603) for partial support for this research.

Christensen, D. 2007. "Does Murphy's law apply in epistemology? Self-doubt and rational ideals." *Oxford Studies in Epistemology* 2: 3–31.

Christensen, D. 2009. "Disagreement as evidence: The epistemology of controversy." *Philosophy Compass* 4: 756–767.

Christensen, D. 2013. "Epistemic modesty defended." In *The Epistemology of Disagreement: New Essays*, edited by D. Christensen & J. Lackey, 76-97. Oxford: Oxford University Press.

Christensen, D. 2014. "Conciliation, uniqueness and rational toxicity." *Nous* 50: 584–603.

Christensen, D. 2016. "Disagreement, drugs, etc.: From accuracy to akrasia." *Episteme* 13: 397–422.

Christensen, D. 2019. "Formulating independence". In *Higher-order Evidence: New Essays*, edited by Mattias Skipper & Asbjørn Steglich-Petersen, 13-34. Oxford, United Kingdom: Oxford University Press.

Christensen, David. 2021. "Akratic (epistemic) modesty." *Philosophical Studies* 178(7): 2191-2214.

Christensen, David. 2022. "Epistemic Akrasia: No Apology Required." *Noûs 1* (online first): 1-22.

Daoust, Marc-Kevin. 2023. "Structural Rationality and the Property of Coherence." *Pacific Philosophical Quarterly* 104(1): 170-194.

Hawthorne, J., Isaacs, Y., & Lasonen-Aarnio, M. 2021. "The Rationality of Epistemic Akrasia". *Philosophical Perspectives* 35(1): 206-228.

Hjortland, O. T. 2017. "Anti-exceptionalism about logic." *Philosophical Studies* 174: 631-658.

Horowitz, S. 2014. "Epistemic akrasia." *Noûs* 48(4): 718-744.

Kolodny, N. 2005. "Why be rational?" *Mind* 114(455): 509-563.

Lee, Wooram. 2024. "What is Structural Rationality?" *Philosophical Quarterly* 74(2): 614-636.

Lin, Hanti. 2023. "Bayesian Epistemology." In *The Stanford Encyclopedia of Philosophy* (Winter 2023 Edition), edited by Edward N. Zalta & Uri Nodelman, URL = <https://plato.stanford.edu/archives/win2023/entries/epistemology-bayesian/>.

Pryor, J. 2018. "The merits of incoherence." *Analytic Philosophy* 59(1): 112-141.

Quine, W. V. O. 1951. "Two Dogmas of Empiricism." *The Philosophical Review.* 60(1): 20-43.

Read, S. 2019. "Anti-Exceptionalism about Logic." *Australasian Journal of Logic* 16(7): 298.

Scanlon, T. M. 2003. "Metaphysics and morals." *Proceedings and Addresses of the American Philosophical Association* 77(2): 7-22.

Williamson, T. 2007. *The Philosophy of Philosophy*. Oxford: Blackwell Publishing

Worsnip, A. 2021. *Fitting Things Together: Coherence and the Demands of Structural Rationality*. Oxford University Press.

Worsnip, Alex. 2022. "Making space for the normativity of coherence." *Noûs* 56(2): 393-415.