Are There Indefeasible Epistemic Rules?

Darren Bradley

University of Leeds

1. Introduction

What if your peers tell you that you should disregard your perceptions? Worse, what if your peers tell you to disregard the testimony of your peers? How should we respond if we get evidence that seems to undermine our epistemic rules? Several philosophers have argued that epistemic rules are indefeasible. I will argue that all epistemic rules are defeasible. The result is a kind of epistemic particularism, according to which there are no simple rules connecting descriptive and normative facts. I will argue that this type of particularism is more plausible in epistemology than in ethics. The result is an unwieldy and possibly infinitely long epistemic rule — an Uber-rule. I will argue that the Uber-rule applies to all agents, but is still defeasible — one may get misleading evidence against it and rationally lower one's credence in it.

Section 2 explains the problem of undermining and three possible responses. Section 3 explains the indefeasibility view, and section 4 argues against it. Section 5 argues against contributory rules, and Section 6 contains the core of the positive proposal, arguing that the problem can be solved by understanding epistemic rules as hedged. Section 7 develops the proposal using a generalization of the concept of admissible evidence. Section 8 extends and defends the resulting position — where the only unhedged rule is a single Uber-rule. Section 9 compares our position to Quinean holism. Section 10 concludes.

2. The Problem of Undermining and Three Responses

In this section I will explain the problem of undermining, and the response that epistemic rules are indefeasible. In the next section I will argue that this response is ad hoc.

Consider some epistemic rules that have been discussed in the recent literature:

Testimony If an agent's situation[1] includes testimony that x, then the agent is rationally required to believe that x. (Compare Elga [2007], Titelbaum [2015].)

Perception If an agent's situation includes a perception that x, then the agent is rationally required to believe that x. (Compare Chisholm [1966], Huemer [2000], Pryor [2000], Boghossian [2008].)[2]

Credence-chance link If an agent's situation includes full belief that the chance of x is y%, then the agent is rationally required to have credence in x of y%. (Compare Lewis's [1980] Principal Principle.)[3]

Let me make some clarificatory comments about these rules. First, these rules have descriptive (by which I will mean non-normative) concepts in the antecedent and normative concepts in the consequent, making them narrow-scope rules.[4]

Second, the supervenience of the normative on the non-normative entails that there will be some such narrow-scope rules. One can take them to be grounding principles.

---

1. The term 'situation' is used by most authors in this literature. Titelbaum offers a definition of 'situation': "I will assume only that whatever the true theory of rationality is, it may specify certain aspects of an agent's circumstances as relevant to determining which overall states are rationally permitted to her. Taken together, these relevant aspects comprise what I'll call the agent's 'situation.'" I take it this is compatible with all authors' usages.

2. It is controversial whether this rule requires a different treatment, as how things look is not a psychological attitude. I think similar rules will still hold for non-attitudinal psychological states, but it won't matter for my arguments.

3. Lewis formulated the Principal Principle in terms of conditional epistemic probabilities e.g. $Cr(A \mid \text{Known chance of } A \text{ is } x \text{ and } E) = x$, and I do not intend to diverge from this version. I have used a conditional to show the similarity to the other rules. More on this in section 7.

4. Narrow-scope rules have the form [if p then you ought to believe q]. A wide-scope rule would have the form 'you ought to believe [if p then q]'. See Broome (1999).

Third, the normative concepts in the consequent should be understood in the epistemic sense, as opposed to the pragmatic or any other sense. For example, if you are offered $1m to believe in God, there is a sense in which you ought to believe in God. I am not concerned with this pragmatic sense of 'should'.

Fourth, the normative concepts in the consequent should be understood as relative to the agent's situation. There is a sense of 'ought' in which you ought to believe p iff p is true.[5] There is also a sense of 'ought' in which you ought to believe p iff p is justified in your situation. It is this latter sense which I will use.[6]

Fifth, I take 'defeasible' to mean 'can be rationally doubted'.

Sixth, I will call rules like these "simple" to distinguish them from the complicated Uber-rule I will defend in sections 6 and 7.

Seventh, the antecedents contain specific positive claims about agents' states, but do not contain totality facts e.g. 'the agent has no other perceptions/testimony'.

Our problem is that it is easy to generate cases where narrow-scope rules conflict or undermine themselves. For an example in ethics, suppose a crazed murderer asks you where their target is. The principle not to lie conflicts with the principle to prevent others from coming to harm.[7] For an example in epistemology, suppose a table looks red but you are told that it is blue. Perception conflicts with Testimony.

The epistemology literature has focussed on cases where a rule is *self-undermining*. Start with a case of peer disagreement:

> Suppose that you and a friend independently evaluate a factual claim, based on the same relevant evidence and arguments. You become confident that the claim is true. But then you find out that your friend — whose judgment you respect — has become just as confident that the claim is false. Should that news at all reduce your confidence in

---

5. See Fassio (2018 2.b.) for discussion of the truth norm of belief.

6. Hedden 2012 p. 344.

7. This is a problem for Kantian ethics. See Timmerman (2013).

the disputed claim? Conciliatory views on disagreement answer "yes." (Elga 2010 p. 175)

Elga (2007) initially defended a conciliatory view. For simplicity, we can consider the most extreme conciliatory view:

> *Testimony* If an agent's situation includes testimony that x, then the agent is rationally required to believe that x.

But Elga found that many of his colleagues disagreed with him. Elga writes:

> Suppose that you have a conciliatory view on disagreement, but you find out that your respected [colleague] disagrees. He has arrived at a competing view (about disagreement), and tells you all about it. If your conciliatory view is correct, you should change your view. You should be pulled part way toward thinking that your friend is right. In other words, your view on disagreement requires you to give up your view on disagreement. (p. 179)

So Testimony looks unstable: if Elga receives testimony against Testimony, and applies Testimony, he must give up Testimony.[8] Conciliatory views sometimes call for their own rejection. This is the self-undermining problem.

There is an important implicit premise in Elga's argument we should make clear – Anti-Akrasia.

> *Anti-Akrasia*: It is never rational to believe [x and I should not believe x]

Without Anti-Akrasia a rational agent could believe testimony (i.e. believe x if someone tells you x) and also believe that one should not believe testimony. The former is the lower-level belief x; the latter is a higher-level belief about beliefs. The agent would follow Testimony

without believing Testimony. So Elga's argument against Testimony implicitly assumes that rational agents are never Akratic.

Although I've no interest in defending anti-Akrasia, I think Anti-Akrasia is plausible, and that we cannot block Elga's argument by simply allowing Akrasia.[9] So I will grant Anti-Akrasia and block Elga's argument for a different reason.

Returning to the main thread, let's distinguish three ways of responding to the self-undermining problem. One might ignore colleagues and maintain a high credence in Testimony.[10] When it comes to your belief in being conciliatory, you should stubbornly ignore your colleagues. Conciliatoriness is indefeasible. The result is that rationality has what Titelbaum calls "fixed points" — these are propositions expressing the rules of rationality, and they are indefeasible. These rules have built-in restrictions to ensure they cannot be undermined, so Titelbaum calls them "restricted rules".

A second response, which I will defend, is the view that you *should* be moved by your colleagues. Testimony might start off with a high prior, but you can get evidence against it and decrease your credence, just like any other belief. (And if you rationally disbelieve that some rule is a requirement of rationality, then it doesn't apply to you.) I will develop this view using the distinction between *hedged* and *unhedged*

---

8. See Elga 2010 p. 181–2 and Titelbaum 2015 p. 271.

9. See Horowitz (2014) for a number of powerful arguments against Akrasia. Elga (2010) doesn't mention Akrasia; Titelbaum (2015) discusses Akrasia in detail, but doesn't engage with Elga's argument. Here is an objection to my view: "Your position is that there are no simple unhedged rules; as Anti-Akrasia is a simple unhedged rule, how can you assume Anti-Akrasia?" Response: On the view I will develop, if Akrasia *is* ever rational, it will be rational only in highly unusual cases, e.g. where you rationally believe that an expert tells you: 'p and you should not believe p'. This is a strange situation (compare Worsnip 2018 p. 24), so we should expect to be in a strange belief state. No such situation arises in the case Elga discusses, so I don't think allowing Akrasia is a good way to block Elga's argument. (Note that the rules I discuss are narrow-scope epistemic rules, while Anti-Akrasia is a wide-scope coherence requirement [Worsnip 2018].) Thanks to a referee for pressing this point.

10. Elga (2010) and Titelbaum (2015) defend this type of approach. Elga writes: "[O]ne should be moved by disagreement about some subject matters, but not about disagreement itself …" (p. 184).

rules.[11] Hedged rules have ceteris paribus clauses, stating situations where the rule fails to apply; one such situation is where you rationally disbelieve that some rule is a requirement of rationality. I will argue that rules of rationality are hedged — one can always get evidence against them being requirements of rationality. The result is that there are no "fixed points", no rules that rational agents should always believe to be requirements of rationality — all are defeasible.

A third response is also worth discussing. There is a different way to weaken Testimony, suggested by Christensen (2010, 2013).[12] Distinguish *absolute* and *contributory* rules. Absolute rules have a consequent that says that you are required to be in some state, or to perform some action; contributory rules have a consequent that says that you have *a reason* to be in some state, or to perform some action. Applied to epistemology, absolute rules say what you should (or should not) believe, whereas contributory rules say what counts in favour of (or against) a belief. Testimony is an absolute rule; but consider a contributory version of Testimony:

> *Contributory Testimony*: If an agent's situation includes testimony that x, then the agent has *a reason* to believe that x.

Paradox is avoided because one can have a reason to believe x and a reason not to believe x. But I will reject contributory rules due to cases of "valence-switching", i.e. where what is usually a reason *for* becomes a reason *against*.

Here is a map of the main positions, the sections in which they are discussed and (tentative) suggestions for where some philosophers might be placed:

	Hedged	**Unhedged**
Absolute	Me, Holton? Väyrynen? Sections 6, 7 and 8	Elga, Lasonen-Aarnio, Titelbaum. Sections 3 and 4
Contributory	End of section 6	Ross, Christensen? Section 5

I argue against contributory rules in section 5. I will defend hedged rules in sections 6 and 7. In the next two sections I will explain and then argue against restricted rules.

3. Restricted Rules

Elga (2010) defends the first option — restriction — arguing that we need to make a modification to Testimony to make it immune from defeat. Titelbaum (2015) builds on this and suggests that epistemic rules have the following form:[13]

> *Restricted Testimony*: If an agent's situation includes testimony that x, the agent is rationally required to believe that x — unless[14] x contradicts [Testimony[15]].
>
> If x contradicts Testimony, then ignore x and continue believing other testimony.[16]

---

11. See Holton 2002; Väyrynen 2006, 2009; Schroeder 2004 section 5.

12. See Dancy (2013 section 1). Christensen's (2010 p. 203–4; 2013 p. 92–3) talk of "ideals" could be understood as talk of unhedged contributory rules. His earlier work (e.g. 2007) suggests that there is a residual bad-making feature of violating ideals/rules, in which case the rules would not be contributory.

13. Titelbaum doesn't actually defend Restricted Testimony. His point is that epistemic rules must have this form, whatever they turn out to be.

14. 'Unless' means 'if not'. Roughly, if there is no contradiction, then believe testimony. Notice this says nothing about what to do if there is a contradiction. That's why we need the next line.

15. The original text says "this rule". I assume 'this rule' refers to Testimony. It won't matter much if it refers to Restricted Testimony. I argue that holding such rules to be indefeasible is ad hoc, whatever their exact content is.

16. I've added "If x contradicts Testimony, then ignore x and continue believing other testimony." Titelbaum isn't explicit about how to respond if x

Restricted Testimony cannot undermine itself. It says we should believe testimony on most topics, but not with regard to the question of whether we should believe testimony. It follows that our credence in Restricted Testimony should stay the same, even in the light of opposing testimony. Thus, Restricted Testimony is indefeasible by testimony.

Titelbaum (2015) defends similar restrictions on all epistemic rules, e.g.:

> *Restricted Perception*: If an agent's situation includes a perception that x, the agent is rationally required to believe that x — unless x contradicts [Perception]. (p. 273)

And not only can rules undermine themselves — they can undermine each other, e.g. you might be told that your perception is unreliable. To block Restricted Testimony from undermining *other* rules, Titelbaum defends:

> *Properly Restricted Testimony*: If an agent's situation includes testimony that x, the agent is rationally … required to believe x — unless x contradicts [a] truth about what rationality requires. (p. 274)

This structure is intended to generalize to all rules, e.g.:

> *Properly Restricted Perception*: If an agent's situation includes a perception that x, the agent is rationally required to believe that x — unless x contradicts a truth about what rationality requires.

Titelbaum suggests that we should have credence 1 in the rules of rationality; they are *indefeasible fixed points*. To use his memorable phrase: mistakes about rationality are mistakes of rationality:

> [E]very agent possesses a priori, propositional justification for true beliefs about the requirements of rationality

_____

contradicts Testimony, but I think he must be committed to this line. I'll leave this implicit when not needed.

in her current situation. An agent can reflect on her situation and come to recognize facts about what that situation rationally requires. Not only does this reflection provide her with justification to believe those facts; that justification is ultimately empirically indefeasible.[17] (p. 276)

But this is *very* hard to believe. Christensen (2013) writes:

> Suppose … that I follow [Properly Restricted Testimony] and remain absolutely confident in its correctness, despite the fact that it's rejected by many epistemologists I respect, and even rate as my superiors in philosophical skill. How should I view my own reasoning on this topic? Should I think that while I'm generally only moderately reliable when I think about philosophy, nevertheless when I think about arguments for general conciliation, and for not being conciliatory about conciliation, I'm especially immune from error? That seems extremely dubious. (p. 89)

Given the difficulty of formulating rational rules, the claim that we should be *certain*, or even highly confident, of what they are, even in the face of opposing arguments, seems to me untenable.

Furthermore, suppose you do hear testimony against some (correct) rule of rationality. Titelbaum suggests that you should not believe such testimony to *any* degree. It naturally follows that you should not

17. Titelbaum adds a footnote, saying that the rules "could be opposed by empirical evidence pointing in the other direction …. But those propositional justifications are ultimately indefeasible in the sense that the empirical considerations will never outweigh them and make it all-things-considered rational for the agent to form false beliefs about what her situation requires." I don't understand this. If empirical considerations can count against the rules, why can't we imagine increasingly strong empirical considerations that eventually outweigh the rules? And the restrictions he places on the rules seem to ensure that empirical considerations can*not* count against rational rules at all. And his later claim (section 6) that the Fixed Point Thesis leads to the Right Reasons view suggests that empirical evidence cannot count against rational rules.

believe you have *any reason*, no matter how weak, to reject that rule of rationality. We are led to the view that *what seems like evidence doesn't even count as evidence*. Lasonen-Aarnio[18] is sympathetic to this view, writing that it "rests merely on a desire to avoid paradox" (p. 342). But this position seems at least as paradoxical as any of the alternatives.

So where did we go wrong? I will argue in the next section that Elga's (2010) initial move to a restriction on Testimony — a move extended by Titelbaum — was a step in the wrong direction.

4. Elga's Argument for Ubiquitous Indefeasibility

When faced with a rule that undermines itself, Elga modifies the rule so that it ignores evidence that threatens to undermine itself. But this looks ad hoc. What reason do we have to believe this, other than the fact that it avoids the undermining problem? After all, most of our beliefs are subject to doubt in the light of opposing evidence, so why are beliefs about disagreement different?

Elga argues that *all* fundamental epistemic rules must be non-underminable. He writes:

> In order to be consistent, a fundamental … rule … must be dogmatic with respect to its own correctness. (p. 185)

But what does 'fundamental' mean here? Elga (2010) writes:

> a fundamental [rule] is one whose application is not governed … by any other [rule] …. (p. 179)

This is close to saying that fundamental rules are those that are not defeasible. But then the question is: Why should we think that there are any such rules? Indeed I will argue that there are none.

Elga argues that fundamental rules are dogmatic using the following example: Imagine a magazine, *Consumer Reports*, consistently rating itself as the best consumer magazine.[19] When faced with the complaint that the magazine is biased when it recommends itself, Elga endorses an editor saying:

> To put forward our recommendations about toasters and cars is to put them forward as good recommendations. And we can't consistently do that while also claiming that contrary recommendations are superior. So our always rating ourselves #1 does not result from an arbitrary or ad hoc exception to our standards. We are forced to rate ourselves #1 in order to be consistent with our other ratings. (p. 185)

But I think this brings out the importance of the distinction between fundamental and non-fundamental rules. The advice of a consumer magazine is not a plausible candidate to be a fundamental rule, and it is easy to imagine cases where *Consumer Reports* allows that its recommendations can be defeated. Suppose the editor of *Consumer Reports* knows that her rival magazine has a larger budget and, as a result, makes recommendations based on more evidence. This editor should recommend that consumers rate this rival magazine's recommendations over those of *Consumer Reports*. It would be ad hoc for the editor to ignore evidence against her magazine's recommendations.

It is only *fundamental* epistemic rules that are plausibly (non-adhocly) indefeasible. Elga needs to argue that *there are* simple fundamental epistemic rules, and, to get the consequences he wants regarding the peer disagreement debate, he needs to argue that Restricted

---

18. "It may come as a surprise that in some cases a state can be perfectly epistemically rational even if one has what would *seem* like strong evidence for thinking that it is not" (Lasonen-Aarnio 2014 p. 342, italics added). Although the next sentence does seem to count the apparent evidence as real evidence. And she then says, "That one should believe that one shouldn't φ doesn't entail that one shouldn't φ" (p. 343). So it seems Lasonen-Aarnio does accept that there is evidence against φ, which affects higher-level beliefs but not first-order beliefs, and thus she accepts Akrasia.

19. Based on Lewis (1971).

Testimony (or something similar) is among them.[20] But no such arguments are offered. So restricting Testimony to make it indefeasible is ad hoc after all.

Someone might object that what is fundamental is some other rule that is more complicated than Testimony (and Titelbaum takes no stand on what the fundamental rules are). But once we start making qualifications, we are on the path to the Uber-rule and the view that all rules are defeasible. This is the path I think we should be on. Before presenting my positive view, I will argue in the next section against contributory rules in epistemology. (Those uninterested in contributory rules can skip the next section without loss of continuity.)

5. Against Contributory Rules

Distinguish two types of defeasible rules — rules can be *outweighed*, or they can be *voided* (and possibly both). Frederick (2015) makes the distinction as follows:

> [N]ormally, if Alf promises Betty that he will meet her for lunch, Alf thereby acquires a duty to meet Betty for lunch. However, if Alf's promise to Betty was made under threat of force, his promise fails to engender that duty, because the circumstance was *duty-voiding*.
>
> Normally, if Alf has the duty to meet Betty for lunch, then he ought to meet Betty for lunch. However, if on his way to meet Betty he sees a child drowning in a pool, and he can save the child without much risk to himself, then Alf has a duty to save the child .... If Alf cannot both save the child and meet Betty for lunch, he has two duties which conflict. Perhaps all would agree that Alf ought to save the child; but his duty to meet Betty remains. In the

> circumstances, his duty to save the child [outweighs][21] his duty to meet Betty. (p. 256)

Let's say that rules which can be outweighed are *contributory*. And we can take it that the distinguishing feature of contributory rules is that there is a residual bad-making feature when they are not followed, even when not following them is the right thing to do (due to other, weightier rules).[22] When rules are voided, there is no residual bad-making feature. In this section I will argue against contributory rules in epistemology.

In meta-ethics, contributory rules are associated with Ross (1930), who argued that we have numerous duties, and what we ought to do depends on the overall weighting of these duties. In epistemology, the analogous view is that we have numerous epistemic reasons to believe, and what we epistemically ought to believe depends on the overall weight of these reasons. We can make this explicit by weakening the consequent of our rules:

> *Testimony*: If an agent's situation includes testimony that x, then the agent is rationally required to believe that x

to

> *Contributory Testimony*: If an agent's situation includes testimony that x, then the agent has *a reason* to believe that x.

(Assume for now that the rule is unhedged.) This avoids the undermining problems above. An agent who has testimony that x and a perception that not x has a reason to believe x and a reason to believe not x.[23] No paradox; what they should believe depends on the correct weighing of these reasons. And an agent who doesn't believe that

---

20. This is worth emphasizing. Although Elga is not explicit that conciliationism is a fundamental rule, it must be one if Elga's argument that fundamental rules are indefeasible is to have implications for conciliationism.

21. Frederick uses 'overridden' but I prefer 'outweighed', which makes explicit that they still have weight.

22. Schaffer 2015 p. 659.

23. Similarly, testimony against Testimony is compatible with Testimony still providing *a reason* to believe testimony.

Contributory Testimony is a requirement of rationality still plausibly has a reason (a weak one) to believe testimony.

My objection is that contributory rules say that a feature that is a reason to believe x is always a reason to believe x. And it is plausible that in some cases a feature that usually is a reason to believe x becomes a reason to believe not x. And it is especially plausible in epistemology with regard to Perception:

> [I]n a case where I … believe that I have recently taken a drug that makes blue things look red and red things look blue, the appearance of a red-looking thing before me is reason for me to believe that there is a blue, not a red, thing before me. It is not as if it is some reason for me to believe that there is something red before me, but that as such a reason it is overwhelmed by contrary reasons. It is no longer any reason at all to believe that there is something red before me; indeed it is a reason for believing the opposite. (Dancy 2013)

These strong *undercutting defeaters*[24] provide an argument for *reasons holism*: "a feature that is a reason in one case may be no reason at all, or an opposite reason, in another" (Dancy 2004 p. 7).

The consequence here is that even the contributory version of Perception is refuted. There are situations in which the perception of a red-looking thing is *no reason at all* to believe it is red. Similar examples for Testimony can be devised, perhaps where the agent finds herself inside a logic puzzle on the island of Liars. So we should reject contributory epistemic principles as too strong.

Schroeder (2011) objects that there is still a reason, just a weaker one.[25] He uses the following case:

24. Undercutting defeaters suggest that one's ground for the belief is not sufficiently indicative of the truth of the belief — the ground here being the experience. See Pollock 1967.

25. Schroeder (2011 fn. 8) does not actually deny reasons holism. Nevertheless, the argument he offers can be naturally understood as an argument against reasons holism.

> In the basic case, you are standing outside the library, when you see Tom Grabit exit, pull a book from under his shirt, cackle gleefully, and scurry off. This gives you pretty good reason to believe that Tom just stole a book from the library.
>
> Case 2 is just the same as the first case, except that Tom has an identical twin, Tim, from whom you can't visually distinguish him. In this case, it has seemed to the judgment of many philosophers that your visual evidence is not a reason to believe that Tom stole a book. (p. 333–4)

But Schroeder thinks your visual evidence remains a reason to believe that Tom stole a book. To demonstrate this, he extends the example:

> Consider a third version of the case, exactly like the other two except that in the third case, in addition to Tim, Tom has a third identical sibling, Tam, from whom you can't visually distinguish him. This third case underwrites a compelling argument against the intuitive judgment that in the second case, your visual evidence was no reason to believe that Tom stole the book. For if you go on to conclude, in the third case, that Tom stole the book, then you are doing worse than if you had gone on to conclude this in the second case. Your reason to believe that Tom stole the book therefore doesn't seem to have gone away in the second case; it merely seems to have gotten substantially weaker. (p. 334)

Schroeder seems to be right about this case, where the defeater weakens the strength of the reason. But I don't see how he can say the same about the previous case in which a drug makes blue things look red and red things look blue. In that case, the defeater changes the valence of the reason, i.e. what was a reason *for* becomes a reason *against*. So contributory rules are still too strong.

Furthermore, a motivation for wanting contributory reason in ethics does not apply to epistemology. A strong motivation for contributory reasons in ethics is a need for outweighed reasons. Dancy writes:

> Scanlon [who is taken to reject contributory rules[26]] has … deprived himself of the idea of a defeated reason, and thereby prevented himself even from addressing the question what the appropriate response is to such a thing. Normally we would speak of regret and residual duties, but if all conflict is, as Scanlon suggests, merely apparent, there are no defeated [outweighed] considerations capable of demanding regret, and nothing to generate a residual duty.[27] (Dancy 2004 p. 26)

The idea that contributory reasons keep their force even when defeated is most plausible in ethics, especially if we think of our ethical system as consisting in duties. ~~Our duties remain even if they are defeated by other duties.~~ Your duty not to lie remains, even if you have a stronger duty to save a life by lying.

By contrast, our *epistemic* duties do not seem to have the same power to remain even if defeated. For example, suppose you see a red-looking table in a situation where you know you have taken a drug that makes blue things look red and red things look blue. You don't believe that the table is red. Do you regret that you have not lived up to your epistemic duty to believe things are as they appear to be? Surely not. A signature feature of a violation of a contributory rule is that it involves a residual bad-making feature. But there doesn't seem to be anything bad about believing the table is blue. So the contributory view looks to be unmotivated in epistemology. So the rejection of contributory rules is more plausible in epistemology than in ethics.

26. I am neutral on whether this is a good interpretation of Scanlon.

27. This argument is central to Dancy's position. He later (p. 28–9) rejects Holton's Principled Particularism for the same reason.

6. Hedged Rules

Let's recap. So far I have argued that it is ad hoc to hold that simple epistemic rules are indefeasible, and that epistemic rules are not contributory. My positive view is: *For any (simple or complex) epistemic rule, rational agents can acquire evidence that it is not a requirement of rationality, causing them to decrease their credence that it is a requirement of rationality.* If the credence is low enough, the rule does not apply to the agent (by Anti-Akrasia).

Let's connect this to familiar views in ethics. Recall Frederick's example: if Alf's promise to Betty was made under threat of force, his promise fails to engender that duty, because the situation was duty-voiding. So, starting with a simple ethical rule:

> if you promised to p, then you are required to p

it is plausible that you have *no reason* to p if the promise was made under duress. The simple ethical rule can be voided. In order to allow for this, we need the full rule to be:

> if you promised to p, then you are required to p, unless the promise was made under duress.

We'll say that the full rule is *hedged*.[28]

In epistemology, using the example of Perception, a better specification of the rules would move us from:

28. This idea has been suggested in meta-ethics by Holton (2002) and Horty (2007). Horty writes that "the general principle that lying is wrong should be taken to mean simply that lying is wrong by default — that is, to a first approximation, that once we learn that an action involves lying, we ought to judge that it is wrong, unless certain complicating factors interfere" (p. 23). Holton suggests that ethical rules need 'That's it' clauses stating that there are no other ethically relevant features, e.g. 'Any action that has such-and-such features and That's It is wrong'. One choice-point here is whether the hedge lists a manageable number of exceptions (e.g. 'under duress') or is a place-holder for an open-ended list of exceptions (e.g. 'and there are no other relevant normative features'). This is the topic of section 8. See Field (2000 p. 135 and Appendix) for related points.

Perception: If an agent's situation includes a perception that x, then the agent is rationally required to believe that x

to

Hedged Perception: If an agent's situation includes a perception that x, then the agent is rationally required to believe that x, unless they have evidence against Perception.[29]

What if an agent does have evidence against Perception? I suggest we take Hedged Perception to include:

If they have evidence against Perception,[30] then their credence that Perception is a requirement of rationality in this situation should decrease.

Where there is moderate evidence that Perception is not a requirement of rationality in this situation, the rational agent might retain a middling credence that Perception is a requirement of rationality in this situation. Where there is overwhelming evidence that Perception is not a requirement of rationality in this situation, the rational agent might disbelieve that Perception is a requirement of rationality in this situation. Assuming Anti-Akrasia, (which links beliefs about the requirements of rationality with requirements of rationality), it follows that Perception is not a requirement of rationality in this situation. A fortiori, Perception is not a requirement of rationality in all situations.

A referee objects that where believing E is based on Perception, which is then undermined by E, the support for E vanishes, thus the rule is believed again, and the position is diachronically unstable. In

response, I suggest that the agent can move straight to a stable state in which E has higher credence than before and that-Perception-is-a-requirement-of-rationality-in-this-situation has lower credence than before.

So situations where rules appear to conflict are in fact situations where the rules are incompletely specified. This avoids the smell of ad hoccery — we are not asserting that rules are immune to conflicting evidence; we are placing limits on when the rule applies, a move which is familiar in ethics and philosophy of science.[31] It is also familiar in epistemology — it is the strategy suggested by Lewis (1980) regarding objective chance which I develop in the next section.

7. Two Types of Inadmissibility

Let's start with a simple credence-chance link:

Credence-Chance Link: If an agent's situation includes full belief that the chance of p is x, then the agent is rationally required to have credence of p in x.

This rule does not apply in all situations. An agent can have evidence that justifies their having credences that differ from the chances. Call such evidence *inadmissible*. The most familiar form is evidence that gives us direct information about the event. For example, if you see a fair coin landing Heads in a crystal ball you know to be reliable, you should not have credence of ½ that it will land Heads. This is an *opposing* defeater. So the Credence-Chance Link should be hedged. Indeed, Lewis defended a hedged version of the Credence-Chance Link — the Principal Principle (PP), of which we'll use a simplified version:[32]

---

29. Notice that although Hedged Perception contains 'unless', it is very different from Titelbaum's restricted rules. Restricted rules tell us to ignore *evidence* that conflicts with the rules; hedged rules tell us to lower credence in the *rules* (in our situation) when there is evidence undermining them.

30. I intend this to cover cases where the evidence is that Perception fails to apply in their current situation, and cases where the evidence is that Perception fails to apply in all situations.

31. Compare Cartwight (1983), Pietroski and Rey (1995).

32. Precisely, it says: "Let C be any reasonable initial credence function. Let t be any time. Let x be any real number in the unit interval. Let X be the proposition that the chance, at time t, of A's holding equals x. Let E be any proposition compatible with X that is admissible at time t. Then C(A∣XE) = x" (p. 266). This brings in various features which aren't relevant to our concerns. One that is worth mentioning is that Lewis's rule requires that agents not only have credences that match the known chances, but also update in such a way

PP If an agent's situation includes full belief that the chance of p is x, then the agent is rationally required to have credence of p in x, *unless they have inadmissible evidence*.

But there is a second type of inadmissible evidence that has not been discussed in the objective chance literature — the Credence-Chance Link can have *undercutting* defeaters. There might be evidence that reduces your confidence that a rule is a requirement of rationality.[33] Suppose an apparently reliable agent tells you that chance is not something which your credences should match. You might be confused by such a statement, but this confusion is surely enough to justify your credences' not perfectly matching the chances. Such evidence reduces confidence in the Credence-Chance Link.

Similarly, suppose an apparently reliable agent tells you that the PP is false (i.e. even with the hedge). Then you should lower your credence that *PP* is a rule of rationality. So PP also needs to be hedged. Just as the hedge of the *Credence-Chance Link* generated the PP, we need a new principle generated by the hedge of the PP.

One way to systematize all this is to make the concept of admissibility relative to a rule. So we can restate PP as:

PP*: If an agent's situation includes full belief that the chance of p is x, then the agent is rationally required to have credence of x in p, *unless they have evidence-inadmissible-relative-to-the-Credence-Chance-Link*.

And PP* needs to be hedged in turn:

Qualified PP*: If an agent's situation includes full belief that the chance of p is x, then the agent is rationally required

---

that, given any possible admissible evidence, they will continue to do so. We could set up all our rules in the same way, but I will use the simpler formulation. Thanks to x.

33. Such evidence most directly reduces your confidence that a rule is a requirement of rationality in your situation; a fortiori it reduces your confidence that a rule is a requirement of rationality in all situations.

---

to have credence of x in p, unless they have evidence-inadmissible-relative-to-the-Credence-Chance-Link,   *or unless they have evidence-inadmissible-relative-to-the-PP**.

I suggest that other rules work the same way, and that we need to generalize the notion of inadmissibility to apply to the other rules.

Let's work through how this applies to Testimony.

Testimony: If an agent's situation includes testimony that x, the agent is rationally required to believe that x.

Like the Credence-Chance Link, this needs to have exceptions built in. One type of exception occurs when the agent has other evidence against x. Perhaps the agent has directly seen that x is false; this is an opposing defeater. A different type of exception would be if the agent rationally believes that the person testifying is unreliable; this is an undercutting defeater. Or suppose an apparently reliable agent tells them that Testimony is false; this is a more general undercutting defeater, relevant to other testifiers too. Call evidence that justifies an agent in not believing testimony *inadmissible-relative-to-Testimony*. So Testimony should be hedged as follows:

Hedged Testimony: If an agent's situation includes testimony that x, then the agent is rationally required to believe that x, *unless they have evidence that is inadmissible-relative-to-Testimony*.

Mutatis mutandis for other simple epistemic rules (including Hedged Testimony). Thus, I suggest that all simple epistemic rules are hedged. The undermining problems disappear, as the problem cases described situations where the agent does have evidence inadmissible to some rule, so the agent has low credence that the rule is a requirement of rationality, so by Anti-Akrasia (which links beliefs about the requirements of rationality with the requirements of rationality), the rule does not apply.[34] We have arrived at my view that epistemic rules will be

---

34. Objection: Our hedged simple rules will almost never apply. Hedged

not the simple rules with which we began, but complex rules full of hedges.

Dancy (2004) attributes an analogous view to Scanlon (1998):[35]

> Scanlon's view … seems to be that … there are no actual conflicts, only appearances of conflict. If two of our [rules] seem to get in each other's way, what this shows is that at least one of them is incompletely specified, and the matter is resolved by a more complete specification. Suppose we [can help someone in need at the cost of killing someone else].[36] The idea here is that in a proper understanding of the [rule] that requires us to help those in need, there would probably be included an exception to that duty for all cases where to help one we have to kill another. Properly understood, therefore, the duty to help cannot conflict with the duty not to kill. (p. 25)

But a new danger emerges if the complexity spirals out of control. Starting with simple rules, can the exceptions be finitely stated? Ideally, we would like to have finite exceptions, as this would allow a manageable set of rules that could be used to guide our deliberation.

I don't know if this is possible, so I will concede the point, and defend the possibility that the exceptions are open-ended. The idea is that, even if we are left with an infinite list of exceptions to simple

---

Perception might start: "If it looks like p, then believe p, unless you have reason to distrust your vision or …." Anyone who has ever experienced an illusion does have some such reason. This point is correct — hedged simple rules might rarely apply to any actual people — but this is not a problem. Again, compare the Principal Principle. Lewis formulated it to apply only to initial credence functions, so it does not apply to any real people. Nevertheless, it is (if correct) a substantive and informative rule of rationality.

35. Dancy cites Scanlon (1998 p.197–200). But Dancy admits that the interpretive claim is not beyond dispute, and I refrain from attributing this to Scanlon.

36. Bizarrely, the original sentence is "Suppose we face a choice between killing one person and helping another." I take it Dancy intended to describe the example I use.

---

epistemic rules, this position is still defensible as an epistemic theory. So in the next section I will defend an Uber-rule.

Before that, there is a loose end to tie up. Suppose we accept hedged rules — are they contributory or absolute? In the ethics case, we used:

> *Hedged, Absolute*: If you promised to p, then you are required to p, unless the promise was made under duress, etc.

But the alternative is:

> *Hedged, Contributory*: If you promised to p, then you have a reason to p, unless the promise was made under duress, etc.

The advantage of contributory rules is that they don't undermine themselves or each other; but we have no need of this advantage here, as the hedge can ensure that there is no conflict. And if there is no conflicting rule, then surely a promise to p means that you are required to p; the Hedged, Absolute rule is correct.

The epistemology case looks similar.

> *Hedged, Absolute*: If it looks red, then you should believe it is red, unless you rationally believe you have recently taken a drug, etc.

> *Hedged, Contributory*: If it looks red, then you have a reason to believe it is red, unless you rationally believe you have recently taken a drug, etc.

The stronger Hedged, Absolute rule looks plausible, so I tentatively endorse it.

8. For the Uber-rule

So far I have argued that the link between descriptive and normative concepts can be described only by a tapestry of interlocking hedged

rules. Can these rules be finitely stated? If not, we have a version of particularism:

> *Principled Particularism*: Any finite set of rules will be in-sufficient to capture all normative truths.[37]

Thus the connection between descriptive and normative truths is expressible only with an infinitely long rule: an *Uber-rule*. With this one Uber-rule, which presumably does not undermine itself, the problems of undermining are avoided.

I will defend the Uber-rule in epistemology. As far as I can tell, the term 'Uber-rule' was introduced by Christensen (2010) and described as follows:

> Suppose we specify, for every possible evidential situation in which an agent may find herself, what the appropriate doxastic response is. The result would be an overarching rule which took into account every sort of evidence. We might then think of that rule as encoding the one and only true epistemic [rule]. (p. 203)

I add that the Uber-rule cannot be finitely expressed. This position can be generated by conjoining an infinite number of simple rules, or positing a finite number of simple rules with at least one infinitely long hedge. As the Uber-rule applies in all situations, it is unhedged; I will argue below (8.2) that it is also defeasible.

So far the only detailed discussions of the Uber-rule have argued against it.[38] Let me sketch a way of thinking about rationality that makes an Uber-rule plausible. Instead of thinking about rationality as emerging out of simple rules, think of the Bayesian approach where agents begin with a prior distribution of probabilities. Imagine all epistemically possible worlds on a vast Venn diagram. Bayesians only allow updating by conditionalization (or Jeffrey conditionalization),

so the prior distribution encodes what agents should believe given any evidence. The prior distribution can encode an Uber-rule.

Are there any simple rules that apply in all situations, i.e. all over the Venn diagram? I see no reason to expect so. We cannot say you should believe your senses, because there is possible evidence / areas of the Venn diagram where you should not believe your senses. Indeed, for any simple rule you might state, there is possible evidence / areas of the Venn diagram where you should not follow that simple rule. The best we could hope for would be patterns in some areas of the Venn diagram that could be helpfully described with simple rules.

We might make a comparison with Humean laws. For the Humean, laws do not explain events; events are fundamental and laws are derived from events, so Humean laws are just informative summaries. Similarly, in Bayesian epistemology the priors are fundamental; some simple epistemic rules might be derived from the priors, but, like Humean laws, they are just helpful summaries.[39]

For a further intuitive argument, consider some difficult question, such as the correct credence that climate change is man-made. Is it plausible that a finite number of simple epistemic rules would generate the rational credence? It strikes me as entirely implausible.

In the rest of this section we'll consider two objections to the Uber-rule based on (i) guidance and (ii) coherence.

8.1. Guidance
Lasonen-Aarnio (2014) offers the most detailed discussion, and

---

37. Compare Holton (2002) and McKeever and Ridge (2006 p. 16).

38. Christensen (2010, 2013), Lasonen-Aarnio (2014).

39. Thus simple rules can still help *explain* rational requirements; we just have to remember that the simple rules are themselves explained by the priors. So this seems compatible with Christensen's comment: "If we ask why the disagreement of other competent thinkers with the same evidence should affect my confidence, the correct *explanation* may still be that since their disagreement is evidence that my initial belief was based on an epistemic error, it creates rational pressure to give credence to the claim that my initial belief was based on error, and that … this creates rational pressure to back off of that initial belief to at least some extent" (Christensen 2013 p. 93, italics added). However, Christensen (2010 p. 203–4) offers some considerations that suggest that some simple rules are explanatory in a stronger sense than I allow.

focusses on two main worries. Start with the worry that we cannot be guided by the Uber-rule:[40]

> Now, the problem for the Uber-rule view is that an Uber-rule just doesn't seem like the kind of rule that can offer genuine guidance. For one, it cannot even be expressed as a set of finite, informative generalisations. … Even if one argues that subjects manage to genuinely follow the Uber-rule by employing more ordinary kinds of epistemic rules as heuristic guides, the fact remains that they need guidance to follow the Uber-rule itself. Hence, the Uber-rule is a very awkward candidate for a rule that is itself supposed to play the role of offering genuine guidance. (p. 333)

In my view the Uber-rule need not offer guidance. I prefer the Bayesian view above, according to which agents are "guided" by their priors and the evidence.

Still, Lasonen-Aarnio's worry may survive as the worry that we mere mortals are unable to be guided by the full ideal prior probability function, which requires assigning probabilities to an infinite number of propositions. So let's address this worry in the form that Lasonen-Aarnio puts it — that the Uber-rule cannot offer guidance.

In response, it is arguable that an Uber-rule could offer guidance. In the happiest cases, agents know that the ceteris paribus clause of a simple rule is true, i.e. other things are equal, so they can be guided by the rule. For example, someone who has conclusive reason to believe that their senses are reliable can follow Perception, and Perception will be one part of the larger Uber-rule. Though never guided by the entire Uber-rule, agents might usually be guided by the part of the rule that is relevant to them — it looks red, there are no other relevant features, so believe it is red.[41]

40. Compare Boghossian 2008 p. 496.

41. See Väyrynen (2008) for a detailed defence of a similar position in metaethics.

However, this position might be harder to defend in epistemology than in ethics. In ethics, often only a small number of considerations are relevant to whether a particular action is right. But, turning to epistemology, consider whether I should trust my vision and believe there is a red table in front of me. Presumably the track-record of my vision is relevant, so every visual experience in my life — and whether it was veridical — will be relevant. And we already seem to have a rule that is too complicated to guide me.

At this point we could retreat to the view that although we can never (or not always) say that there are *no* other relevant features, these other relevant features can be rationally ignored. Defending the view that there are more reasons than we normally take into account, Mark Schroeder (2005) writes:

> If God made a list of all of the pros and cons of some course of action, it might be infinitely long. But you can't possibly take everything into account — only the reasons near the top of the list. (p. 15)

So we could be guided by the Uber-rule by paying attention to the most important parts of the rule for our situation.

A problem might remain if someone insisted that guidance must be provided by the *full* epistemic theory, i.e. the whole Uber-rule. The full Uber-rule cannot be finitely formulated, and perhaps we cannot be guided by rules we cannot formulate.

But an agent might be guided by the rules without being able to formulate them. For example, when you judge that a sentence is ungrammatical, you are guided by linguistic rules that you are unable to formulate.

Furthermore, you might even be disposed to get the rules wrong if you tried to formulate them. For example, Arpaly (2003) argues that Huckleberry Finn is guided by the true moral rules in not turning in an escaped slave, even if he wrongly believes that he should turn in

the slave.[42] Huckleberry Finn's inability to formulate the rules does not stop him from being guided by the rules.

Perhaps our position in epistemology is analogous to Huckleberry Finn's position in ethics. We cannot formulate the Uber-rule any more than Huckleberry Finn can formulate the ethical rules. But it is possible that his compassion makes him perform the right action, so there is a sense in which he is guided by ethical rules. Similarly, it is possible that our good sense, or epistemic intuition, makes us form rational beliefs, so there is a sense in which we are guided by the Uber-rule.[43]

To sum up, although the Uber-rule is incompatible with some views that involve strong requirements on our ability to formulate and be guided by normative rules, such strong requirements can be rejected.

8.2 Coherence

Lasonen-Aarnio's other worry is that there might be no Uber-rule to be found — "finding a rule not susceptible to defeat is surely harder than merely defining one to be such!" (p. 331). The idea seems to be that the concept of an Uber-rule might be incoherent, just as the concept of a square circle is.

One might immediately be suspicious here — surely, for any possible evidential state, there is a rational response, and the Uber-rule states what the rational responses are. Nevertheless, Lasonen-Aarnio offers the following case to defend this position:

> [A]ssume that you are staring at a chart representing the Uber-rule: for each possible epistemic situation … the chart specifies what the recommendations made by the

Uber-rule in that situation are. … Now imagine that you hear an epistemology oracle tell you that the recommendations made by the Uber-rule in the very situation you are in right now are incorrect. In so far as the rule is complete … the chart must say something about your current situation. Imagine that, as the chart tells you, the rule recommends being in state S. But in so far as the oracle is to be trusted, doesn't her testimony act as a higher-order defeater for any such recommendation? (p. 331)

As stated, I think this imagined situation is incoherent. The problem is that we are imagining a case where you are told that the Uber-rule is incorrect. But we have defined 'Uber-rule' as the *correct* rule. So any rule that fails to be correct isn't the Uber-rule. So any speaker who tells you the Uber-rule is incorrect is saying something incoherent, and you should not follow their advice.[44] Even worse, we've been told that it is an "epistemology oracle" who is telling us the Uber-rule is incorrect. Presumably, the phrase 'epistemology oracle' applies only to someone who speaks the truth. So 'epistemology oracle' cannot be correctly applied to anyone who tells us the Uber-rule is incorrect (under any mode of presentation), as the Uber-rule is by definition correct.

A coherent scenario in this area is that a rational agent might be told something false about the Uber-rule by a plausible-looking but misleading informant. And what might be worrying Lasonen-Aarnio is the thought that, in this scenario, agents can have rational doubts about the content of the Uber-rule. It looks paradoxical for the Uber-rule to apply in all situations, be unhedged, and yet be defeasible.[45]

I think paradox can be avoided. Let's work through a case. Suppose the sum total of your evidence is a current red experience. Suppose the

---

42. See Raz (2000) for further discussion. One might be tempted to appeal to the familiar distinction between a theory of rightness and a decision-making procedure (Bales 1971). But we are working only with the subjective ought, so the theory of rightness looks irrelevant (or perhaps better: the distinction collapses).

43. A different problem with an infinitely long Uber-rule is that we would be unable to grasp epistemic concepts. Thus, Jackson, Pettit and Smith (2000) give a semantic objection to particularism. This seems to be a good reason for positing a long but finite rule.

44. They might say, "The rule you are following is incorrect", but that's a different situation, as it involves a different mode of presentation of the Uber-rule.

45. Compare Christensen (2013): "If the agent continues to follow the Uber-rule while doubting its correctness, it seems inevitable that she will in some cases violate [Anti-Akrasia]" (p. 93). The example shows how we can doubt the Uber-rule without Akrasia.

Uber-rule says that the rational response to this evidence is to be 90% certain that there is a red object in front of you. You are rational, so you are 90% certain that there is a red object in front of you. And you are reflective, so you have the second-order belief that the Uber-rule says that agents with the sum total of a red experience should be 90% certain that there is a red object in front of them. Say you are n% certain of this second-order belief. Now add an informant who says (falsely) that the rational response to your current red experience is to be only 50% certain that there is a red object in front of you.

Distinguish first-order and second-order responses. The most immediate response is lowering your *second*-order credence that the Uber-rule says that agents with exactly the evidence of a red experience should be 90% certain that there is a red object in front of them. This falls from n%. The *first*-order response is that your credence that there is a red object in front of you drops below 90%. If you fully trusted the informant, it would fall to 50%, but let's say it settles at 70%.

There is no violation of the Uber-rule. Your epistemic position has changed — your evidence now includes the red experience *and the testimony*. And it is compatible with this story that the recommendation of the Uber-rule for someone with this evidence is "Have 70% credence that there is a red object in front of you."

No paradox so far. Perhaps the problem is that you can doubt the content of the Uber-rule. You are not certain that the Uber-rule recommends a 90% credence that there is a red object to agents with just a red experience. Similarly, you should not be certain that 70% is the rational credence in your current situation. But a rational agent may doubt the Uber-rule without believing that the beliefs it prescribes are irrational. The possibility that your credence should be more than 70% needs to *perfectly balance* the possibility that your credence should be less than 70%.[46]

46. What about Titelbaum's memorable phrase — mistakes about rationality are mistakes of rationality? I can endorse this phrase if we fill it out as follows:
Titelbaum's-Principle-I-Accept: For any given agent and situation, if they are mistaken about what they should believe in their situation, then they are making a mistake of rationality.

As a final point in defence of the Uber-rule, it is worth emphasizing Holton's point that Principled Particularism is compatible with utilitarianism:

> [S]uppose you were a utilitarian. Then you couldn't determine which action to perform on the basis of a list of pleasures and pains caused by possible actions of yours. You would need to know, in addition, that these were all the pleasures and pains that each action caused; and that these were all your possible actions. (p. 206)

No matter how much of the world we describe, it is not enough to ensure that an act is right — rightness depends on the whole world, just as generalizations like 'all swans are white' do.[47] Similarly, if we focus on narrow-scope rules with descriptive antecedents without totality facts, then we can never stop at finite rules — what you ought to believe depends on your whole mental state. But this doesn't rule out

This follows from the internalist intuition that agents have epistemic access to the rationality-makers. But one can have rational false beliefs about what others should believe:
Titelbaum's-Principle-I-Deny: For any given agent and situation, if they are mistaken about what one should believe in a situation other than their own, then they are making a mistake of rationality.
The difference can be put in terms of the scope of the quantifiers. For any agent and situation, there are epistemic rules they ought to believe are requirements of rationality; but there are no epistemic rules they ought to believe are requirements of rationality in all situations. Another way of putting this is that you can be rationally mistaken about what someone else should believe, but not about what you should believe. Why the asymmetry? Because what an agent should believe depends on *their own* higher-order beliefs about rationality, whereas what an agent should believe does not depend on *someone else's* higher-order beliefs about rationality. The asymmetry can be read off the Anti-Akratic rule (this is an ad hominem criticism of Titelbaum, as Anti-Akrasia is central to his theory) — it is irrational to believe p and believe that it is irrational to believe p; it is not irrational to believe p and believe that it is irrational for someone else to believe p. Titelbaum (2015) mentions this possible asymmetry but immediately rejects it: "[E]very plausible story I've been able to come up with is generalizable: it applies just as well to an agent's conclusions about what's rationally required in other situations as it does to conclusions about what's required in her current situation" (p. 276).

47. Compare Schroeder 2011.

generalizations. So this version of particularism is not devastating for normative theorizing.

This completes my defence of the Uber-rule. In the next section, I will connect the resulting position to broader issues regarding defeasibility and the a priori.

9. Against Certainty

I argued in the previous section that even the Uber-rule can be rationally doubted. This supports a view associated with Quine — that no statement is immune from revision. Epistemic rules are good candidates for statements that are immune from revision, so by arguing that they are not immune from revision, the general case that no statement is immune from revision is supported.

Still, I have only discussed narrow-scope rules connecting descriptive with normative statements. My position does not *entail* that no statement is immune from revision. For example, one might still hold that rational agents are certain of tautologies.

Nevertheless, I think rational doubt can be raised even about tautologies, and for similar reasons. For example, suppose a heavenly voice tells you that p-and-not-p, or tells you that your credences should sum to 0.8. Baffling situations for sure, but they seem to provide some reason to doubt tautologies.[48]

This is a challenge to Bayesianism, which models agents as probabilistic, and so requires that agents are certain of tautologies. I think the Bayesian should respond that their models, like most models, make idealizations — that is, they make assumptions that are known to be false in order to make the model easier to work with. Probabilism is such an assumption. So probabilistic agents in Bayesian models are *idealized*; but they are not necessarily *ideal*, in the sense of being perfectly rational. We have unfortunate terminology where 'ideal' can be

used to refer to either of these properties. (Compare: There is nothing rational about an ideal gas. The assumption of the idealness of a gas is analogous to the assumption of the probabilism of agents.[49])

In one sense, this type of revisability is stronger than Quine allowed, for one way Quine thinks we might give up a *sentence* is a way that we would naturally describe as a change in the meaning of the sentence. (Quine denied that we could separate meaning change from belief change, due to his rejection of the analytic/synthetic distinction.) Whereas I am happy to talk about *propositions* and inclined to hold that no proposition is immune from revision. But there are related Quinean views that I do not endorse.

First, I do not reject the a priori. I think the Principal Principle (and other principles) are a priori yet defeasible. Some philosophers identify the a priori with immunity from revision — so I reject the "a priori" only in this sense.[50]

Second, I do not reject analyticity. Let an analytic sentence be a sentence that one can be in a position to justifiably believe in virtue of understanding it. This allows the existence of defeaters that block the justification. One can reject an analytic sentence as false if one mistakenly doubts that it is analytic.[51] Imagine being told that 'All bachelors are men' is false by a misleading but eminent source. Or being told that you have ingested a drug that generates false beliefs about which sentences are analytic. It might be rational to reject the sentence, so even analytic sentences are not immune from revision.[52]

---

48. See Williams (forthcoming) for a defence of "rational illogicality"; Schechter (2013) offers related arguments against being certain. At odds with this is the literature based on McFetridge (1990) (e.g. Hale 2002, Ahmed 2000, Leech 2015), which seems to assume that there must be some rules which are immune to doubt.

49. Hájek 2006.

50. Casullo (2003) argues that there is no experiential indefeasibility condition in the concept of a priori justification; and Summerfield (1991) and Thurow (2006) argue that a priori justification is defeasible by experience.

51. Alternatively, perhaps analytic sentences can be rejected in the sense that they are believed to fail to usefully apply to the world. For example, one can reject the sentence 'Ether conducts heat' on the grounds that there is no ether. Eklund (2017 p. 89) holds that analytic sentences can be false.

52. Williamson (2007) writes that "the central idea behind epistemological conceptions of analyticity is that, in such cases, failure to assent is not merely good evidence of failure to understand; it is constitutive of such failure" (p. 73). I reject the analytic in this sense.

10. Conclusion

The problems of this paper were generated by the paradox that emerges in situations where epistemic rules undermine themselves or each other. One way to avoid paradox is to maintain that epistemic rules are indefeasible and ignore all opposing evidence. I have argued instead that we should think of simple epistemic rules as hedged rules. They apply only if agents don't have evidence that is inadmissible relative to those rules. I have defended the view that the only rule that applies in all situations is an Uber-rule which states what agents should believe given any possible evidence. But even the content of the Uber-rule can be rationally doubted.[53]

References

Ahmed, A. (2000). Hale on Some Arguments for the Necessity of Necessity. *Mind* 109 (433):81–91.

Arpaly, N. (2002). *Unprincipled Virtue: An Inquiry Into Moral Agency*. Oxford University Press.

Bales, R. E. (1971). Act-Utilitarianism: Account of Right-Making Characteristics or Decision-Making Procedure? *American Philosophical Quarterly* 8 (3):257–65.

Boghossian, P. A. (2008). Epistemic Rules. *The Journal of Philosophy* 105 (9):472–500.

Broome, J. (1999). Normative Requirements. *Ratio* 12 (4):389–419.

Cartwright, N. (1983). *How the Laws of Physics Lie*. Oxford University Press.

Casullo, A. (2003). *A Priori Justification*. Oxford University Press.

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

——— (2010). Higher-Order Evidence. *Philosophy and Phenomenological Research* 81 (1):185–215.

——— (2011). Disagreement, Question-Begging, and Epistemic Self-Criticism. *Philosophers' Imprint* 11 (6):1–22.

——— (2013). Epistemic Modesty Defended. In Jennifer Lackey and David Christensen (eds.), *The Epistemology of Disagreement: New Essays*. Oxford University Press.

Dancy, J. (2004). *Ethics without Principles*. Clarendon Press.

——— (2007). Defending the Right. *Journal of Moral Philosophy* 4 (1):85–98.

——— (2013). Moral Particularism. *The Stanford Encyclopedia of Philosophy* (Fall 2013 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2013/entries/moral-particularism/>.

Eklund, M. (2017). *Choosing Normative Concepts*. Oxford University Press.

Elga, A. (2007). Reflection and Disagreement. *Noûs* 41 (3):478–502.

——— (2010). How to Disagree About How to Disagree. In Ted A. Warfield and Richard Feldman (eds.), *Disagreement*. Oxford University Press.

Field, H. (2000). Apriority as an Evaluative Notion. In Paul Boghossian and Christopher Peacocke (eds.), *New Essays on the A Priori*. Oxford University Press.

Hájek, A. (2006). Masses of Formal Philosophy 'Interview'. http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.124.5339&rep=rep1&type=pdf.

Hale, B. (2002). Basic Logical Knowledge. *Royal Institute of Philosophy Supplements* 51:279–304.

Hedden, B. (2012). Options and the Subjective Ought. *Philosophical Studies* 158 (2):343–60.

Holton, R. (2002). Principles and Particularisms. *Proceedings of the Aristotelian Society* 76 (1):191–209.

Hooker, B. and Little, M. (eds.) (2000). *Moral Particularism*. Oxford University Press.

Horowitz, S. (2014). Epistemic Akrasia. *Noûs* 48 (4):718–44.

Horty, J. (2007). Reasons as Defaults. *Philosophers' Imprint* 7(3):1–28.

Huemer, M. (2000). Direct Realism and the Brain-in-a-Vat Argument. *Philosophy and Phenomenological Research* 61 (2):397–413.

Jackson, F., Pettit, P., and Smith, M. (2000). Ethical Particularism and Patterns, in Hooker and Little 2000.

Lackey, J. and Christensen, D. (eds.) (2013). *The Epistemology of Disagreement: New Essays*. Oxford University Press.

Lasonen-Aarnio, M. (2014). Higher-Order Evidence and the Limits of Defeat. *Philosophy and Phenomenological Research* 88 (2):314–45.

Lewis, D. (1971). Immodest Inductive Methods. *Philosophy of Science* 38 (1):54–63.

––– (1980). A Subjectivist's Guide to Objective Chance. In Richard C. Jeffrey (ed.), *Studies in Inductive Logic and Probability*. University of California Press.

McFetridge, I. G. (1990). Logical Necessity: Some Issues. In John Haldane and Richard Scruton (eds.), *Logical Necessity & Other Essays*, volume 11 of Aristotelian Society Series. Aristotelian Society.

McKeever, S. and Ridge, M. (2006). *Principled Ethics: Generalism as a Regulative Ideal*. Oxford University Press.

Pietroski, P. M. and Rey, G. (1995). When Other Things Aren't Equal: Saving Ceteris Paribus Laws from Vacuity. *The British Journal for the Philosophy of Science* 46 (1):81–110.

Pollock, J. L. (1967). Criteria and Our Knowledge of the Material World. *The Philosophical Review* 76 (1):28–60.

Pryor, J. (2000). The Skeptic and the Dogmatist. *Noûs* 34 (4):517–49.

Ross, W. D. (1930). *The Right and the Good*. Oxford University Press.

Schaffer, J. (2015). What Not to Multiply without Necessity. *Australasian Journal of Philosophy* 93 (4):644–64.

Schechter, J. (2013). Rational Self-Doubt and the Failure of Closure. *Philosophical Studies* 163 (2):429–52.

Schroeder, M. (2004). The Scope of Instrumental Reason. *Philosophical Perspectives* 18 (1):337–64.

––– (2005). Realism and Reduction: The Quest for Robustness. *Philosophers' Imprint* 5 (1):1–18.

––– (2011). Holism, Weight, and Undercutting. *Noûs* 45 (2):328–44.

––– (2013). Scope for Rational Autonomy. *Philosophical Issues* 23 (1):297–310.

Strahovnik, V., Potrč, M., and Lance, M. N. (eds.) (2008). *Challenging Moral Particularism*. Routledge.

Summerfield, D. M. (1991). Modest A Priori Knowledge. *Philosophy and Phenomenological Research* 51 (1):39–66.

Thurow, J. (2006). Experientially Defeasible A Priori Justification. *The Philosophical Quarterly* 56 (225):596–602.

Timmermann, J. (2013). Kantian Dilemmas? Moral Conflict in Kant's Ethical Theory. *Archiv für Geschichte der Philosophie* 95 (1):36–64.

Titelbaum, M. G. (2015). Rationality's Fixed Point (Or: In Defense of Right Reason). *Oxford Studies in Epistemology* 5:253–94.

Väyrynen, P. (2006). Moral Generalism: Enjoy in Moderation. *Ethics* 116 (4):707–41.

––– (2008). Usable Moral Principles, in Strahovnik et al. 2008.

––– (2009). A Theory of Hedged Moral Principles. *Oxford Studies in Metaethics* 4:91–132.

Williams, R. (forthcoming). Rational Illogicality. *Australasian Journal of Philosophy* 96 (1):127–41.

Williamson, T. (2007). *The Philosophy of Philosophy*. Blackwell Publishing.

Worsnip, A. (2015). The Conflict of Evidence and Coherence. *Philosophical and Phenomenological Research* 96 (1):3–44. doi:10.1111/phpr.12246.