

In *Accuracy and the Laws of Credence*, Richard Pettigrew compiles and expands on his recent work as part of a popular movement in formal epistemology, which Pettigrew calls “veritism”. Veritists hold that *accuracy* is a central source – or for Pettigrew, the only source – of epistemic value. Building on that, Pettigrew (along with others engaged in “accuracy-first epistemology” or “epistemic utility theory”) takes a decision-theoretic approach to epistemic rationality, starting from veritism and deriving rational requirements. The aim of this book is to show that several plausible norms of epistemic rationality can be justified using this approach. Here Pettigrew argues that veritism is quite powerful: it can yield probabilism, a version of the Principal Principle, the Principle of Indifference, and “Plan Conditionalization” (a synchronic version of conditionalization).

One of the major contributions of this book is Pettigrew’s work in articulating mathematically precise arguments for the rational norms he defends, spelling out which versions of these principles can be defended on the basis of which additional assumptions. I defer completely to Pettigrew on all mathematical matters, so I won’t get into the details of these arguments. Instead, I will focus my comments here on some bigger-picture thoughts about the argument for the Principle of Indifference (“PoI”), raising some worries both about its underlying assumptions and about the version of the PoI that it aims to establish. Before doing that, I will make a broader observation about the overall structure of the book.

In the introduction to the book, Pettigrew bills his project as a partial argument for veritism. By showing that veritism can lead to certain plausible conclusions (such as probabilism), Pettigrew aims to lend support to the view.

However, a closer look at the arguments reveals that they do not straightforwardly build on one another: we can’t necessarily accept all of veritism’s potential good consequences at the same time. That’s because while Pettigrew’s arguments all proceed from a shared premise about epistemic value, they differ in which *decision rule* they use. The arguments for coherence and the Principal Principle rely on various dominance norms. The arguments for the Principle of Indifference (“PoI”) relies on Maximin; later,

Pettigrew uses the Hurwicz criterion to characterize different attitudes towards epistemic risk (in contrast with the extreme risk aversion that leads to the PoI).¹ One of his arguments for Plan Conditionalization relies on Expected-Utility Maximization. In all, the book discusses 18 different decision rules, most of them variations on Dominance.

Some of these decision rules can be used together, but others conflict. Though Pettigrew does not devote much time to looking at which rules conflict and which ones do not, he does discuss one instance of tension in his discussion of the PoI and the Principal Principle.² The PoI is justified using Maximin, and the Principal Principle is justified using a chance-dominance. Each principle is established independently, with detailed arguments in favor of each decision rule. However, the two decision rules are straightforwardly incompatible: they make different recommendations in response to the same set of choices.

At first glance this seems to be a case in which veritism can't do everything we ask of it, at least not all at once. But Pettigrew has a different response: he introduces yet another decision rule which yields both the PoI and the Principal Principle. This principle, "C-Maximin", gets the job done, but it is quite unintuitive. (It says to choose the credal option whose *worst-case expected accuracy*, as assessed by possible chance functions, is best. The rule therefore encodes a particular way of assessing worst-case scenarios as well as a particular way of deferring to chances.³) Should we accept C-Maximin? The only argument for the principle is that it justifies good results. What we are left with is two independent arguments for the PoI and the Principal Principle, and a third, comparatively weaker argument for both. Pettigrew's project therefore leaves genuine room for debate regarding which norms veritism can support *together*, and what we should do with conflicts between different decision rules.

Attending to this aspect of Pettigrew's project helps clarify how we should use the book. Instead of one single sustained picture of rationality, Pettigrew offers a Choose Your Own Adventure: decide on a decision rule, and the book will tell you where it leads.

¹ See Pettigrew [2014b], pp. 168-171. All references are to Pettigrew [2014b] unless otherwise noted. The Hurwicz criterion is a decision rule that yields a recommendation based on a fixed weighting of the best-case and worst-case outcomes.

² p. 173-177.

³ p. 173-177.

⁴ The term "superbaby" comes from Hájek, who credits it to Lewis. Pettigrew does not use this term here, but does in his [2014a], the paper on which some of the PoI chapter is based.

If readers want to see how many adventures they can reasonably choose at once, there is more work to be done. (Even if two decision rules are compatible, one might not want to endorse both.) This book is just the beginning.

We can see that much of the action lies in figuring out which decision rules are acceptable if we turn to Pettigrew's main argument for the Principle of Indifference. I will summarize the argument, raise some worries about its underlying assumptions, and then compare it to one of the alternative arguments for the PoI that Pettigrew discusses.

Pettigrew's favored argument for the PoI (the "Argument from Accuracy", pp. 159-165) goes roughly as follows. Suppose you're a "superbaby", or an agent at the very beginning of your epistemic life.⁴ You have no evidence and you have no credences yet, and you must choose an initial credence function. The credences you choose will turn out to be more or less accurate, depending on how the world is. In these circumstances, Pettigrew argues, it is not reasonable to take a bigger risk than absolutely necessary: if your credences are riskier than others you could have had (i.e. if their worst possible accuracy score is worse than another set of credences' worst possible accuracy score), it is not reasonable to have those credences. In other words, superbabies ought to choose their initial credences using Maximin. That means that superbabies will follow the PoI – specifically, the version of the PoI that Pettigrew favors. This principle says that when an agent's credences range across propositions that form a partition, she should (initially) assign equal credence to each element of the partition. Pettigrew intends this principle to apply to the *richest* partition that the agent's credences range over (propositions are as finely grained as they can be). This avoids paradoxes like van Fraassen's cube factory.⁵

In Pettigrew's argument, Maximin has a special role in that it is only meant to be used once. Once you are no longer a superbaby, you can go ahead and use different decision rules that might tolerate more risk. Here Pettigrew draws a comparison to Rawls, who famously uses Maximin in a similar way. Rawls asks us to imagine agents behind the veil of ignorance, choosing the initial setup of society without knowing (or having

⁴ The term "superbaby" comes from Hájek, who credits it to Lewis. Pettigrew does not use this term here, but does in his [2014a], the paper on which some of the PoI chapter is based.

⁵ p. 167.

any evidence regarding) what their position in that society will be. As far as each agent knows, she could end up in that worst-off situation: so, Rawls argues, she should not allow any greater risk than she absolutely must. Agents behind the veil should use Maximin to choose the setup of society. But after that initial choice, they may reasonably abandon Maximin and follow a less extreme decision rule.

Pettigrew's argument for the PoI relies on three substantial assumptions. First, it assumes that we can determine which initial credences are reasonable by looking at the hypothetical choices made by superbabies. Second, it assumes that it might be reasonable for superbabies to use a different decision rule from the rest of us. And finally, it (explicitly, as a premise) assumes that Maximin is a reasonable rule for superbabies to use. I have doubts about all three assumptions. Let's look at them in reverse order.

First: why Maximin? Pettigrew writes that while he finds Maximin plausible, he can't give it much further defense: it is "normative bedrock".⁶ He then quotes William James, writing that our attractions or aversions to doxastic risk are "only expressions of our passional life" – a remark that doesn't suggest normative bedrock so much as arbitrary preference. I find this unsatisfying, especially since Maximin yields such unreasonable recommendations under normal circumstances. A maximally risk-averse believer would never form any opinions unless she could be sure she was right; a maximally risk-averse actor would never leave the house. So can we give any further argument for why superbabies should follow Maximin?

Rawls' reliance on Maximin has been criticized for the same reason. (Harsanyi ([1975], p. 595), arguing against Rawls, writes that "[i]f anybody really acted this way he would soon end up in a mental institution.") But looking to Rawls' responses in favor of Maximin does not provide much help. In his [1974], where he explicitly addresses criticisms of Maximin, Rawls raises several practical and psychological benefits of the rule. He also writes that the principle is only meant for large-scale, one-time decisions, and not for everyday life.⁷ The first line of thought is clearly unacceptable in the epistemic context, especially given the aims of Pettigrew's project, and the second is unconvincing. As Harsanyi points out, there is no principled distinction between "macro"

⁶ P. 166

⁷ Rawls [1974], p. 142

and “micro” choices – and in any case, it doesn’t seem like one’s choice of decision rule should depend on scale.

Another point from Rawls does seem applicable here: since Maximin only looks at the value of possible outcomes and not their likelihood, it can be applied without any empirical information, or indeed without any credences at all.⁸ This is a clear benefit of Maximin: superbabies, who have no credences, could follow it. But this argument for Maximin is inconclusive. It doesn’t distinguish Maximin from other (incompatible) rules like Maximax.

Let’s turn now to other assumptions of Pettigrew's argument, and the appeal to superbabies in general. Why should superbabies use a special decision rule? Without a significant backstory, it seems arbitrary to treat an agent’s *very first* decision differently from others. Rawls, of course, has an elaborate contractualist justification for the significance of decisions made behind the veil of ignorance: for him, it is not arbitrary to single out this very first choice. But Rawls’ story does not carry over to the epistemic case.

And in both the moral/political and epistemic cases, it seems odd to switch decision rules midstream. Compare two agents who use different decision rules: Ariel maximizes expected utility, and Belle follows Maximin. Ariel should clearly view Belle as unnecessarily cautious, and Belle should view Ariel as reckless. Neither would be rational to allow the other to make decisions on her behalf. So why should someone like Ariel authorize someone like Belle to choose her priors, thereby setting the stage for the rest of her epistemic life?⁹ To put it differently, suppose Ariel found out that her own

⁸ Ibid., p. 143

⁹ In *A Theory of Justice*, Rawls does argue that it is rational for agents to use Maximin behind the veil partly *because* they are making decisions on behalf of others. Ironically, Rawls is making this argument *against* the Principle of Indifference, which would allow agents behind the veil to assign probabilities and then choose the setup of society by maximizing average utility. He writes:

Now I shall assume that the parties discount likelihoods arrived at solely on the basis of [the PoI]. This supposition is plausible in view of the fundamental importance of the original agreement and the desire to have one’s decision appear responsible to one’s descendants who will be affected by it. We are more reluctant to take great risks for [our descendants] than for ourselves.

(Rawls [1999], p. 146)

past, superbaby self was as risk-averse as Belle. If Pettigrew is right, how should Ariel view this past self? Perhaps she should view her own past self as irrational. But if she's been rational all along, shouldn't it be possible that she could know this? Alternatively, perhaps she should view her past self as rational, despite the fact that she views other agents who use Maximin as overcautious. Either way, the result is puzzling; changing decision rules partway through one's life leads to a kind of incoherence. An agent's decision rule encodes how she compares outcomes, and what she cares about, when she is making choices. If someone starts off caring very much about avoiding the worst outcome, why should she suddenly start caring less about that once she gets some information?

Finally, let's turn to Pettigrew's first assumption. Why we should defer to superbabies – agents with no *credences* – as opposed to regular babies, who simply have no *evidence*?¹⁰ It is hard to even imagine an agent who is somehow devoid of beliefs, yet able to make rational decisions. One might even think that epistemic rationality is fundamentally a matter of figuring out what the world is like on the basis of one's standing beliefs and inference rules. That picture suggests that constraints on initial credence functions, if there are any, will not come from some pre-credal rational decisions, but from somewhere else: perhaps, from facts about what is supported by a completely empty body of evidence.

I am not yet convinced of the significance of superbabies, or the rationality of Maximin. So I am not yet convinced by Pettigrew's argument for the PoI. However, I am somewhat optimistic about a different argument for the PoI, one which Pettigrew rejects. This is the "Argument from Evidential Support", from White [2009], which says that

But even granting that Rawls is right about our willingness to take risks on others' behalf, in this case we are talking about one agent at different times.

The comparison with Rawls here raises an interesting question: *should* we draw a significant distinction between an agent's decisions for her future self and her decisions on behalf of others? Pettigrew, in this book, is not an advocate of diachronic epistemic norms (he endorses Plan Conditionalization, not Conditionalization) – so maybe he would say no. Still, I doubt that he would want to endorse Rawls' argument in the epistemic case: the considerations that Rawls is alluding to, such as duties to one's descendants, seem thoroughly moral and practical.

¹⁰ The same could be said about Rawls's argument. But while Rawls could argue that it is impossible to rationally assign credences on the basis of no evidence, Pettigrew can't make that move: he is arguing that superbabies *can* rationally assign credences, using Maximin.

when one's evidence is "symmetrical", one's credences should be as well.¹¹ I'll end by briefly looking at two related advantages of this argument. First, it does not commit us to the particular version of the PoI that Pettigrew defends, which seems to me to be too strong. Second, it in fact does not commit us to *any* particular initial credence distribution, which makes it compatible with other substantive constraints on priors beyond indifference.

Recall that according to Pettigrew's preferred version of the PoI, one's initial credences should be spread evenly over the richest partition of propositions that one considers. Here is his main motivating example for the principle:

Kazuo knows nothing about the colour of the handkerchief in my pocket. He is more confident that it is yellow than that it is red, and more confident that it is red than that it is blue.¹²

Kazuo's credences are distributed over the partition $\{Yellow, Red, Blue\}$, and he has no information. Yet, he is more confident of some propositions than of others. Kazuo seems irrational, and the PoI can explain why.

However, Pettigrew's PoI runs into trouble in other cases. Consider someone else in Kazuo's evidential situation, who considers this set of propositions: $\{Light\ Blue, Dark\ Blue, Yellow, Red\}$. According to Pettigrew's version of the PoI, this agent should assign $\frac{1}{4}$ credence to each of these possibilities. But seems like the wrong answer: *Light Blue* and *Dark Blue* should not *each* be assigned as much credence as *Yellow* or *Red*. If anything, this agent should be indifferent over $\{Blue, Yellow, Red\}$, and then also divide her credence in *Blue* evenly between *Light Blue* and *Dark Blue*.¹³

The Argument from Evidential Support has the resources to capture the intuitively correct verdicts in both cases. Plausibly, an empty body evidence provides symmetrical support to the elements of the partition $\{Yellow, Red, Blue\}$, but not to the elements of $\{Light\ Blue, Dark\ Blue, Yellow, Red\}$. This is why both Kazuo and his more discriminating counterpart should be indifferent over $\{Yellow, Blue, Red\}$.

¹¹ See pp. 155-156.

¹² This is taken from p 1, with color options simplified for clarity.

¹³ Pettigrew might disagree with me about whether this result is counterintuitive. In his [2014a] he endorses the view that an agent who considers $\{Light\ Blue, Dark\ Blue, Red\}$ should assign equal credence to each (p. 27-28). In the book he also endorses the view that an agent who considers $\{Blue, \sim Blue\}$ should assign credence $\frac{1}{2}$ to each (p. 167).

In fact, the Argument from Evidential Support does not commit us to *any* particular initial credence distribution, or to any general principle about how one's credences should be distributed based on the partition of propositions that one considers. This feature of White's defense of the PoI is part of the point of his paper. White argues that it is "clearly misguided" to see the PoI as "taking purely structural features of a space of possibilities as giving conditions of rational credence." Instead, White writes that the right version of the PoI should place these constraints based on "facts about the balance of reasons."¹⁴ Though White raises this point in discussing supposed counterexamples to the PoI, like van Fraassen's cube factory, it can also explain what has gone wrong with Kazuo's counterpart.

Because it is so noncommittal, the Argument from Evidential Support also leaves us free to accept other substantive rational constraints on priors, such as anti-skepticism or green-rather-than-grue projection. We could turn out to sometimes have "default" reason to believe some elements of a partition more strongly than the others; it could be that our evidence is symmetrical in the handkerchief case, by asymmetrical when it comes to trusting perception. These substantive constraints are controversial, of course, but Pettigrew's Argument from Accuracy seems to rule them out from the start.¹⁵ It is a virtue of the Argument from Evidential Support that it leaves these issues open for debate. Of course, the Argument from Evidential Support is left with a significant challenge in saying *when*, exactly, our reasons are symmetrical, and why Kazuo should be indifferent but his counterpart should not. But maybe this is how it should be: epistemology is hard.

To sum up: I have serious reservations about Pettigrew's defense of the PoI. It rests on dubious assumptions, and it leads to a very strong version of the PoI that is likely false. However, Pettigrew's argument – and the book more generally – provides a valuable service even for the unconvinced. In spelling out the precise ways in which

¹⁴ White [2009], p. 168.

¹⁵ To deny this one would have to argue, for example, that there are *more* real-world hypotheses than skeptical ones, or that a rational agent *should* carve up the real-world hypotheses more finely than the skeptical hypotheses. But this doesn't seem right. If skepticism is irrational, it's not because there are fewer skeptical hypotheses out there, but because they are less plausible.

these rational principles can be justified on veritistic grounds, it opens up new possibilities for debate: instead of focusing on the norms themselves, we can look at the decision rules that make them rational to accept. While Pettigrew does not exactly demonstrate the *power* of veritism (as I noted above, it might be undesirable or even impossible to accept all of his arguments at once), he does demonstrate its *flexibility*. He shows us which principles veritists can establish, and at what price.

Bibliography

Pettigrew [2014b] *Accuracy and the Laws of Credence*. New York: Oxford University Press.

– [2014a] “Accuracy, Risk, and the Principle of Indifference”. *Philosophy and Phenomenological Research*.

Harsanyi [1975] “Can The Maximin Principle Serve as a Basis for Morality? A Critique of John Rawls’s Theory”. *The American Political Science Review*, Vol. 69, No. 2. pp. 594-606.

Rawls, John. [1999] *A Theory of Justice: Revised Edition*. Belknap Press. Rev Sub edition.

– [1974] “Some Reasons for the Maximin Criterion” *The American Economic Review*, Vol. 64, No. 2, Papers and Proceedings of the Eighty-sixth Annual Meeting of the American Economic Association, pp. 141-146.

White [2009] “Evidential Symmetry and Mushy Credence.” *Oxford Studies in Epistemology* 3, 161-86.