Twitter Thread by Eli Tyre

Eli Tyre
@EpistemicHope

I think because the reward structure of being a bio-ethicist rewards saying level-headed sounding, cautious-sounding, conventional wisdom?

Though I'm not sure why that is.

@JeffLadish I guess if you want to radically improve the world, you mostly don't go into a field that is about opining on other people's work, you go into something like Engineering and do the work?

@JeffLadish I note that Nick Bostrom is what a Bioethicist should be: he thinks hard about tradeoffs and risks, and crystalizes concepts like the Unilateralist's curse and black ball technologies.

But this is starting from a place of "The world could be vastly better. How do we get there?"

@JeffLadish I don't know why bioethics doesn't look more like a field of (less smart) Bostroms.

@JeffLadish I guess because everyone's disgust reactions are triggered by the actually good proposals, and in order to make those proposals the consensus of the field a lot of people have to bite the bullet and stick their neck out saying "I know this sounds crazy / absurd / vaguely evil...

@JeffLadish ...to our intuitions / etc, but it is actually the right thing to do."

And there are not enough people who are up for that, to reach a consensus?

@JeffLadish But economics, as a field, _does_ have this property. Economists are famously fine with policies that are abhorrent to the untrained intuition, but that they are confident are actually better on net. That's a good chunk of what economics is.

@JeffLadish But maybe economics has a solid enough theoretical underpinning, that it can manage to be that kind of field, while, the state of consensus in philosophy, and ethics in particular is confused?

@JeffLadish Like, if you have people still fighting over deontology vs. utilitarianism, it is pretty hard to build to a field wide consensus that we should obviously do covid variolation, because it will save lives on net.

@JeffLadish In which case the problem traces back to "Philosophy has bad feedback loops: you don't get to clearly know that you got the right answer. Which means philosophers are incentivized to make interesting and counterintuitive arguments, instead of to steer towards the actual truth."

@JeffLadish Ok. I think this is the answer to the original question: by the standards of science and engineering, you can settle a question to a sufficient degree of precision, and then move on to higher level questions using your answer to the first question as a foundation.

@JeffLadish But by the standards of philosophy, you can practically never settle a question. There is always space for more counter argument. If you settled the question, you wouldn't be able to argue about it any more!

@JeffLadish Which means that you don't build a foundation of answers to basic questions that you can assume, in trying to answer more complicated questions.

To do it right, bioethics would have to, at least in many areas, _assume_ utilitarianism.

@JeffLadish But if you try to assume utilitarianism, a bunch of philosophers will jump on you with many counterarguments and paradoxes and bullets to bite, all of which erode the ability of a bioethics field, as a whole, to stick to its guns about ideas that are counter-intuitive.

@JeffLadish One thing to note here is that the philosophical standards also challenge the foundations of science and engineering. Philosophy as a whole holds that we don't have rock solid reasons to trust our "obvious" answers to questions of epistemology and metaphysics.

@JeffLadish (Like whether there is an external world or whether knowledge is possible.)

@JeffLadish But the scientist and engineers just ignore the philosophers and assume those things anyway, at least for the purposes of doing their science.

And this works. It isn't necessarily philosophically grounded, but it works.

@JeffLadish And the difference is, I think, that the goal of science is to actually land on our all things considered most-correct answer (in part, so that engineering can build cool things like spaceships), while the goal of philosophy is to have air-tight ARGUMENTATION for our answer.

@JeffLadish If you're mostly trying to build spaceships, you don't really care about gettier problems, unless they're fucking up your ability to build working spaceships.

@JeffLadish And if you you just want to know how the sun sines, you don't really care that epistemology isn't grounded, because while you might be confused about the finer points, the finer points of epistemology are not going to get in the way

of your getting correct beliefs about the sun.

@JeffLadish It turns out that you can generally build spaceships, and (more controversially) end up with correct beliefs about the sun, without having a solid irrefutable proof about the nature of "beliefs."

@JeffLadish Now it does turn out that having correct beliefs about what beliefs are, is pretty useful for getting more correct beliefs.

But having correct beliefs about what beliefs are turns out to not be the same thing as having solid irrefutable arguments for your belief about beliefs.

@JeffLadish This, by the way, is my main answer to the objection that I am projecting onto some folks that I talked to about critical rationality, recently.

(To be clear, I DON'T think that I definitely understood the points they made, and I may be responding to a straw-man.)

@JeffLadish But, I said recently that I thought that Bayes is the foundation of epistemology.

https://t.co/TfCQgRqUtv

> Culturally, I am clearly part of the "Yudkowsky cluster". And as near as I can tell, Bayes is actually the true foundation of epistemology.
>
> But my personal PRACTICE is much closer to the sort of thing the critical rationalists talk about (assuming I'm understanding them).
>
> — Eli Tyre (@EpistemicHope) December 13, 2020

@JeffLadish Different people gave different objections to that, but one objection was (if I understand it correctly) "But, where do your probabilities come from? Either you're infinitely certain of them, or you have an infinite stack of probabilities about the probability below them."

@JeffLadish https://t.co/8joOxuxaPp

> What crit rats can\u2019t get past re bayesianism is what then justifies the specific probabilities? Isn\u2019t it an infinite regress?
>
> I have an 80% confidence in X. Does one then have a 100% confidence in the 80% estimate?
>
> — Cam Peters (@campeters4) December 14, 2020

@JeffLadish I think I could give some more detailed arguments about how this works, but I also want to dispute frame of the point being argued (if I understand it) a bit.

@JeffLadish It is totally possible to have a functioning brain / epistemology, that actually works for producing knowledge and steering through the world, which does not justify itself on its own terms.

@JeffLadish Like the question of "Does a Bayesian learning mechanism, work, in practice" is a separate question from "How do we justify that it works?"

@JeffLadish In general, I think many questions of the "but how do we justify it?" stripe, just don't matter very much.

For many things, we can't prove that it works, but it does work, and it is more interesting to move on to more advanced problems by assuming some things we can't prove.

@JeffLadish That is not to say that ALL questions of justification are irrelevant. Most of the time, even, it is very practically important to ask "how do we / can we know this is true?"

@JeffLadish But my goal in asking that is to figure out what's true, to the best of my ability, and to my own satisfaction, not to have an airtight argument that something is true.

@JeffLadish ...
Getting back to the original question, I think my answer was incomplete, and part of what is happening here is some self-selection regarding who becomes a bio-ethicist that I don't understand in detail.

Basically, I imagine that they tend towards conventional-mindedness.

@JeffLadish @threadreaderapp unroll