

Open to Debate: Debate about Effective Altruism

July 10, 2024

Effective Altruism's failures are so great it should be abandoned. Effective Altruism's advocates are right that our world contains desperate misery and injustice and that people seeking to engage positively face hard questions about how to do so. But we don't get guidance from this troubled project.

Effective Altruism's fans describe it as a reason- and evidence-based guide to doing good. But its notions of 'reason' and 'evidence' are not commonsensical. Effective Altruists appeal to randomized controlled trials—or similar—in recommending minimally costly programs they think can be shown, with metrics of welfare economics, to have maximum benefits. That's a highly technical program, not a transliteration of ordinary notions of 'reason' and 'evidence'.

Here's a list of four concerns I have about Effective Altruism:

1. First, human lives are more complicated than Effective Altruists tell us. Before Effective Altruism existed, aid experts knew that relying on randomized trials or similar doesn't guarantee aid interventions that work in real life. And, as a philosopher at Stanford, Leif Wenar, has noted, some interventions that Effective Altruists long touted as doing "the most good" might be more harmful than beneficial.
2. Second, Effective Altruists mostly neglect or dismiss as "ineffective" the coordinated efforts of justice movements to change harmful social structures. But Effective Altruists lack the resources for this negative assessment, and, in making it, they neglect the political roots of misery and injustice and so fail to address the reproduction of misery and injustice.
3. Third, because Effective Altruism rarely recommends social change, it reinforces the *status quo*. This partly explains its success in attracting funds. It now directs the giving of hundreds of millions of dollars annually, records pledges in the tens of billions, and influences US politics. Like some other mega-philanthropies, it has an anti-democratic bent, spending vast sums and shaping policy with little accountability.

Effective Altruism started with an important public focus on extreme poverty and animal suffering. My fourth concern is about how it lost this focus.

4. Prominent Effective Altruists are now calling themselves longtermists who contemplate the prospective wellbeing of trillions of humans they imagine living in the distant future and conclude that combatting risks to human extinction (which they call "X-risks") is so important it's a moral priority that might justify

almost any means, however terrible. This reasoning is frightening, as philosopher Émile Torres has stressed. It's like 20th century dictators taking envisioned future utopias to justify unfathomable atrocities. And now longtermists' funders are rushing to build artificial general intelligence *while* insisting it's a key X-risk. Longtermists are funding programs in AI safety to protect against their funders' *own* AI systems.