

## About this doc:

This document is set up for well-willing EA's to collaborate on a review and discussion of *The Good it Promises, The Harm it Does*.

## Method

1. You can click on the chapters to read them separately. Please do not share freely on the internet.
2. You can add comments, note and context in the separate google docs per chapter, and write a short reflection behind the chapter heading in this doc.
3. Feel free to restructure the doc as you want

## Norms:

- Don't strawman. Stress what parts of the critique are valid and present the author's case in the best possible way. Of course, any factual errors should be corrected.
- Never use hostile or scornful language.
- Rapoport's rules: [Link](#). In short, aim to improve the critique, only then try to critique the critique.

## Note:

Some of the papers are kinda hostile, and reading them can be pretty frustrating. Keep in mind these are heavy topics, this is just one person's account and stop reading if you feel too uncomfortable 😊

## Open question:

Is this book bad-faith criticism? From reading, it is clear that no editor went over it to check if the claims were accurate and represented the thing critiqued well. Also, though not mandatory, checking your criticism with some EA before publishing it, would have been a show of good faith. No current EA's were contacted to respond.

# Contents

## [Foreword](#)

*Amia Srinivasan*

Comments: **Need to shorten**

*MK*: “For most Effective Altruists, these movements are, at best, examples of ineffective attempts to do good; negative examples from which to prescind or correct, not political formations from which to learn, with which to create coalition, or to join. This fact reveals, on the part of Effective Altruism, not only a certain moral and intellectual presumptuousness.”

I found this passage striking, as most of the EA’s I know actually take a lot of inspiration from past liberation and social justice movements, as a show of moral circle expansion. Perhaps we can **provide some evidence** for this. It’s remarkable that without any sources, this disrespect for social movements is presented as a **fact**. What I could concede is that perhaps a deep understanding and respect of the history of social progress could be improved in the movement, to be brought on about the same level as the appreciation for the industrial revolution.

“This reply doesn’t always hit its mark, for some of Effective Altruism’s critics have offered powerful arguments that target the core principles and background moral epistemology of Effective Altruism—have offered, that is, “in principle” critiques of Effective Altruism”. **No sources are offered that link to this powerful criticism.**

“They raise worries about the overwhelming whiteness, middle-classness, and maleness of the Effective Altruist community that many of its members are likely to think irrelevant to the assessment of Effective Altruism’s value”

In Europe, I haven’t met EAs who would think our lack of diversity is **irrelevant**. I wonder where this vibe check is coming from. It is fair to say that in this community, the lived experiences of marginalized groups are not seen as *fundamentally necessary* to practice effective altruism, but I’ve mainly met EA’s who feel awkward with this “whiteness and maleness”.

— JWS notes —

I

*Political critique does not, and should not, merely address what social and political movements say about themselves. Political critique does, and should, also think about what social and political movements do: what effects they systematically bring about in the world, which structures they tend to reinforce, and which people they empower and which they silence.*

I think this kind of critique is Srinivasan’s strongest line of attack, and it continues in the following paragraphs. For example, EAs might say that SBF committing fraud doesn’t give us any reason to update on the question of the expanding moral circle, or the ITN framework, or other foundational EA beliefs. Srinivasan is clearly interested in critiquing

EA politically, as it seeks to grow and gain power/influence to then implement effective ideas and interventions. I believe she would say that what matters is the actual impact EA has on matters of suffering, liberation, and justice, and not whether our favoured moral philosophy or decision theory is perfectly consistent, or 'true' in a Platonic sense. This point of view is further evidenced in this quote:

*The historically most significant social movements—none of which, unlike Effective Altruism, were born out of the activism of Oxford philosophers—have offered complex analyses of the interrelations between different structures of oppression*

## II

*There is every possibility, then, that Effective Altruists will ignore what these voices have to say—or fail to take the time to understand what their significance might be. That would be a deep shame, and what's more, a betrayal of what I believe is a real commitment, on the part of many Effective Altruists, to bring about a better world.*

She states right here that she thinks many EAs are genuinely committed about making the world better. I think this is worth highlighting and keeping in mind, even amongst the rest of a hostile book.

— — —

## Introduction

*Carol J. Adams, Alice Crary, and Lori Gruen*

*Comments:*

*MK: “In the face of a host of criticisms, advocates of EA adopt strategies of response that are both slippery and sticky. ... In this way, we can see EA as an ideology in the insidious sense, a system of belief and practice that covers up systemic injustices embedded in the fabric of existing capitalist societies in a manner that clears the way for the perpetuation of significant wrongs and harms.”*

It is not clear to me how EA is an ideology that covers up systemic injustices? This is the main thesis of the book, that EA clears the way for harms and wrongs, so I hope we can understand this together.

What I can concede is that the EA way to deal with criticisms can feel a bit like “embrace and extinguish”, but perhaps someone with more experience with red-teaming can show some positive and negative examples of how core EA orgs have dealt with criticism.

— JWS notes —

I

*One was that EA’s insistence on its favored quantitative metrics pressures activists to work within the market-structured institutions that in many cases are responsible for the very wrongs they seek to address, pushing them to pursue reformist strategies that contribute to the persistence of harmful institutions. <and following, longer paragraph>*

I think that there are a few distinct criticisms going on in this section:

1. EA measuring things wrong (saying programs are not effective, when they are effective)
2. EA using the wrong metrics (e.g. it doesn’t take into account considerations of justice)
3. EA undervalues the potential impact/tractability of systemic, non-reformist change (an ongoing debate in the EA community I think) <- This includes the critique that EAs reformist streak leads back systems to stay in place, counterfactually causing harm
4. Poor internal EA dynamics damage its ability to identify effective opportunities and support/fund them (e.g. elitism, lack of communication, hostile environments, cases of sexual assault/abuses of power dynamics)

II

*Protest as they may, Effective Altruists get their characteristic orientation from much older utilitarian ideas, reworking them and attracting new attention to them.*

I think this is pretty fair tbf, I think one of Toby’s original ideas for the movement was ‘effective utilitarianism’? I think making EA more value-pluralistic is an interesting philosophical space that might be underdeveloped. Having said that, I think one can still think effectiveness concerns have moral importance in other moral approaches.

III

*EA doesn’t have resources for fundamentally criticizing the pertinent capitalistic structures, and instead tends to speak for preserving them and working within them, in*

*effect reproducing their harms. This collection's contributors challenge this stance, pushing back against structures that contribute to the exploitation of women and Indigenous and racialized people who are made to do care work as well as to the devastation of more-than-human nature.*

I think this is a pretty big crux. Is 'capitalism' a net good or a net harm? Is it necessarily exploitative? Can one separate the benefits of the global economic system from the colonialism used to construct it? This could, and probably should, be a book topic on its own.

— — —

Chapters:

## [1.How Effective Altruism Fails Community-Based Activism](#)

*Brenda Sanders*

*Comments:* “But when I finally took the time to contemplate what was at the root of my resentment, ... by someone who understands how activism works in Black communities.”

I really can't begin to interpret these lines. When I read it, I interpret it as: Black people don't/can't do cost-effectiveness, which makes me very angry. Can someone help me with a better interpretation?

It seems that this author is writing off their resentment for being rejected by a fund-manager.

“As far as the effectiveness of community-based vegan activism, there may never be a magical equation for quantifying the “number of animals saved” based on the number of people who attend a vegan festival, class, or food tasting, but since there's a personal component to our engagement, community activists are able to gauge much more precisely the effect we're having on communities than activists engaging in drive-by activism. By tracking people's progression through our classes, workshops, festivals, and other programming, we can get a clear picture of the impact of our activism.”

MK: So I think the picture they are painting of an EA is someone who wants an equation to predict the impact of your project. That's quite something different from demanding a measurement of success. I think there's value on gauging community attitudes as a proxy of your results, but it must be said that people can **say** they are going plant-based, but still **buy** the same amount of animal products.

I think promoting 'measurement' as something fundamentally white and imperial is extremely degrading and harmful.

The best version of this critique is probably somewhat like: EA's demanding BOTEK's and tractability estimates before funding some projects can lead them to miss out on important non-linear social effects, such as 'hype', 'attitude', 'paradigm shifts' etc, as they are harder to predict.

## 2. Effective Altruism's Unsuspecting Twenty-First-Century Colonialism

Simone de Lima

Comments:

— JWS notes —

I

*First, the organizations tend to operate with a top-down, rather than collaborative, approach in their “branch” countries. Goals and standards for success are increasingly established by the organizations’ central HQs, removed from an understanding of local sociocultural, historical, and political factors.*

I think this can be a very valid area of critique for EA - if we're focusing only on what can be quantified and computed, we're just committing the drunkard's fallacy. I think EA orgs should be considering those local factors in their estimates for how effective specific interventions might be. She also provides a unique case of the critique of reformism:

*A case in point involves the ubiquitous cage-free campaigns across the world led by international organizations.*

Though it isn't clear to me whether EA organisations have been involved in that.

II

*Another ubiquitous, depoliticizing, and homogenizing trend has been to associate veganism with what is deemed by Effective Altruists as the palatable, nonthreatening prototype of human: the lifestyle influencer.*

There's no evidence for this right? This just seems to be straight up false. Reading on I think she's referring to the implicit/revealed preference from EA comms, but even there I think this both incorrect and needlessly hostile.

There's also specific criticism of an EA-related organisation in Brazil, though I can't read the website, so not sure how to judge the specific claim, but again it's about which route to veganism is the most likely one to succeed. It doesn't seem related to the earlier criticisms of EA in Brazil.

III

*With the dysfunctionality of capitalism exposed like an oozing wound, if there is one thing that is made clear it is that the local, small-scale, personal connections that make up community support networks—precisely those that the Effective Altruism playbook looks down on as “ineffective”—are what are supplying literal lifelines for communities.*

This piece was written during the height of the COVID pandemic, and refers to the dysfunctional response in Brazil driven by Bolsonaro's government. But I think the response is that EA should be accounting for these groups. Once again, the question of capitalism comes up again.

— — —

### 3. Anti-Blackness and the Effective Altruist

*Christopher Sebastian*

*Comments:*

*MK:* This paper is written very hastily by someone who seems to have little background with EA's methodology. It is therefore hard to comment on it. The author is an aid-skeptic. He seems to suggest that donating to GiveWell top charities is actively harmful because the donors are largely white. Perhaps someone can help write down if and how AMF, for example, works together with local communities and how the acceptance of aid is weighed into GiveWell's metrics.

A little thing about the "whiteness" of EA.

On several occasions in the book, the "whiteness" of EA is mentioned. I think it is important to reflect on what is being said. Leaving aside how strange it is to try to gain any argumentative weight from the phrase "mind-blowingly white", for me, who has never accepted coupling the term 'white' with 'bad vibes', I think there's some deeper point to make.

Many social justice movements want Well Off People Who Haven't Endured Hardship (WOPWHEH's) to STFU. EA's, generally, want WOPWHEH's to Step The F\*\*\* Up. This is a movement **built** to make people acknowledge their privilege and stop bathing in their wealth, and go **do something**. Our target audience is the exact group of privileged people who this author thinks is overrepresented. I think, together with the UK/US origin of the movement, it is completely within reasonable expectation that EA scores poor on diversity. Also, this is a movement that relies on people's personal ethics. We don't aim to be authoritative or evangelical or mostly even normative, so diversity is less of an issue. I think no EA would blame any member of a marginalized group for not joining EA. However, purely for community epistemics and for its probably good consequences (there's a lot of talent in the communities are underrepresented in EA), I think it is good to make a target out of diversifying the movement. I don't think we should feel all too bad for not attracting more POC's, however, because there's only a certain amount of people willing to join movements, and I find it perfectly logical that POC's willing to join a movement prefer those focussed around working on justice, rather than a movement like EA.



#### [4. Animal Advocacy's Stockholm Syndrome](#)

*Andrew deCoriolis, Aaron S. Gross, Joseph Tuminello, Steve J. Gross, and Jennifer Channin*

*Comments:*

[5. Who Counts? Effective Altruism and the Problem of Numbers in the History of American Wildlife Conservation](#)

*Michael D. Wise*

*Comments:*

[6.Diversifying Effective Altruism's Long Shots in Animal Advocacy: An Invitation to Prioritize Black Vegans, Higher Education, and Religious Communities](#)

*Matthew C. Halteman*

*Comments:*

## [7.A Christian Critique of the Effective Altruism Approach to Animal Philanthropy](#)

*David L. Clough*

*Comments:*

## [8. Queer Eye on the EA Guys](#)

*patrice jones*

*Comments: (CK)*

The whole article focuses on one strawmanned aspect of EA: ‘advocates trying to convince people to go vegan, claiming that this is the most effective way of saving animals.’ Almost all of her claims can be countered by counterexamples of other groups within EA.

Morally, her main claim can be summarized with the following quote:

*“Finally, even when we do succeed in reducing worldwide production of a particular animal product, we’re not so much saving lives as we are preventing animals from being born into lives of captivity, suffering, and slaughter.”*

I guess this can be simply a difference in moral framework. I (and many people involved in EA) would likely claim that prevented future suffering has nearly the same moral weight as reduced present harm, which I consider to be more coherent than jones’s claim.

## [9.A Feminist Ethics of Care Critique of Effective Altruism](#)

*Carol J. Adams*

*Comments:*

## [10.The Empty Promises of Cultured Meat](#)

*Elan Abrell*

*Comments:*

## 11.How “Alternative Proteins” Create a Private Solution to a Public Problem

*Michele Simon*

*Comments:*

## [12.The Power of Love to Transform Animal Lives: The Deception of Animal Quantification](#)

*Krista Hiddema*

*Comments:*



[13. Our Partners, The Animals: Reflections from a Farmed Animal Sanctuary](#)

*Kathy Steven*

*Comments:*

## 14. The Wisdom Gained from Animals Who Self-Liberate

*Rachel McCrystal*

*Comments: CK*

The chapter discusses Fred, a goat who escaped a slaughterhouse and came back to knock down the door and caused other animals to escape. The story is so strange and there are so few clear arguments that it almost reads like satire, but I will try to distill the points being made. The author, a sanctuary owner, argues that Fred (and with him other animals who “self-liberate”, ie escape from captivity) is feared by animal exploiters.

It is unclear why giving more publicity to self-liberating animals would be beneficial for animal wellbeing. This is never explained by the author, but based on the words that are used and references that are being made, one can only assume that the story is shoe-horned into a story of power struggle, in which the rebels (the animals) are kept quiet by those in power (the owners of the slaughterhouse). I base this on

*“When a cow bound for the kill floor runs or a pig fights back when she’s being moved from gestation crate to farrowing crate, those acts are flames of liberation that farmers know would burn down the entire system. So I know why animal agriculture and those who profit from it don’t celebrate freedom and rebellion.”*

One could argue that escaping animals show that they live in terror and pain in slaughterhouses, which hopefully would cause people to free them. However, this is a point the author is specifically *not* making, as she focuses on the cases in which “*monkeys release others from the lab*” and “*elephants trample fences to free their friends*”.

The author is willing to “*drain her entire bank account*” and would have saved Fred even if it “*would cost tens of thousands of dollars to care for him over his life*”, because “*every individual deserves the dignity of liberation*”. She states herself that she does not want to spend the same money to spare the lives of “*some other goats*”, which to me seems to be because those goats are anonymous. Why their individuality does not count, is unclear. The author seems to prefer strong and heroic animals, as if they should set an example to other animals. Ironically, this seems to be a typical case of anthropomorphization, while the author argues for a more animal-centered approach.

EA is barely mentioned. The only mention concerns one EA-backed nonprofit whose leader has never been in an animal sanctuary, which the author considers to show that legislators are detached from the beings they are ruling over. Warning for bureaucrats who have never gotten their hands dirty seems a valid argument to me; however, the chapter does not provide any arguments that this is the case for those involved in EA, or at least more than for those not involved in EA.

I will finish with a quote from the article which summarizes it well:

*I want an animal rights movement where Fred is the hero. Where Fred's bravery and act of liberation is told and celebrated. And where if we were able to save Fred from those who wished him harm, from those who knew his power, our funders and allies would see that his life would be worth it.*

## [15. Effective Altruism and the Reified Mind](#)

*John Sanbonmatsu*

*Comments:*

## 16. Against “Effective Altruism”

Alice Crary

Comments:

— JWS notes — <I found it hard to stay charitable toward the end here>

I

I think this is a republication of a previous critical work Crary has released, but it might be slightly updated for this volume. In general, this work is more general and philosophical rather than grounded in the specifics of the Animal Welfare movement.

*In light of the undisputed impact of EA, and its success in attracting idealistic young people, it is important to forcefully make the case that it owes its success primarily not to the (questionable) value of its moral theory but to its compatibility with political and economic institutions responsible for some of the very harms it addresses. The sincere dedication of many individual adherents notwithstanding, EA is a straightforward example of moral corruption.*

I have a few thoughts here. 1) Why people join EA and the sociological growth of the movement is an empirical question - some research here should have the answer! But Crary doesn't argue for it. 2) It's not clear that calling a movement 'straightforward moral corruption' is a sign you're arguing in good faith. It seems to denigrate people in the very same way they accuse EA approaches of doing. I think it merits a robust response from EA if I'm being honest.

II

*A leitmotif of the institutional critique of EA is that this bias is politically dangerous because it obscures the structural, political roots of global misery, thereby contributing to its reproduction by weakening existing political mechanisms for positive social change.*

This is a well-written and succinct summary to be fair, regardless of disagreement on the issue.

*The result is that, despite its sincere pro-animal stance, Animal Charity Evaluators is at risk of strengthening an industrial agricultural system that reproduces horrific animal suffering on a massive scale.*

Can these risk not be quantified? Or at least a structured argument for it presented? One thing I find difficult about this essay is that just assumes this is true/EA is wrong, instead of arguing for it more directly.

III

*Effective Altruists invite us to regard the rightness of a social intervention as a function of its consequences, with the outcome involving the best states of affairs counting as doing the most good. This strategy appears morally confused when considered in terms of the ethical stance of the philosophical critique. To adopt this stance is to see the weave of the world as endowed with values that reveal themselves only to a developed sensibility. To see things this way is to make room for an intuitively appealing conception of actions as right insofar as they exhibit just sensitivity to the worldly circumstances in question.*

Maybe this is just going to a fundamental crux? There's not much progress to be made here to be honest, apart from stating the positions more clearly. But while there's room to compromise on the institutional questions, there's very little to be done here. (Crary does point this out later, but again doesn't argue for it - furthermore, her criticism really takes down a whole raft of philosophical traditions and worldly ideologies, of which EA would only be one victim).

#### IV

*This means that the tendency of EA-affiliated organizations to wrongly prioritize evaluation of the proximate effects of particular actions is not a fixable methodological flaw. The organizations focus on these evaluations because it is only here that their image of the moral enterprise seems plausible. It is often right to act in ways that aim to improve the welfare of others. But recognizing the instances in which this is (or isn't) right requires capacities for engaged social thought that EA disavows.*

I'm not sure I understand this argument to be honest, but from what I do understand it doesn't seem to logically flow from the points made earlier in the piece? Once again, a lot of statements but few arguments.

*These interrelated literatures supply additional backing for the verdict that EA, with its abstract methods, bars itself from dealing responsibly in social assessments.*

Only if you accept these methods. Part of EAs tradition might just be rejecting these as philosophically wrong, even if they can help in specific cases.

*In practice, the composite critique suggests that, within any domain in which they operate, charities guided by EA ratings will in general direct funds toward simple interventions capturable with metrics such as income levels or health outcomes, and in a manner relatively insensitive to whether these interventions contribute to perpetuating the institutions that reliably produce the ills they address, while also disparaging as less "effective" systematic attempts to change these institutions.*

I just literally think that this is not argued for in the article at all. It's stated, but never argued for, and has little to do with the philosophical debate. This is still an empirical argument, right? Where EA-backed interventions are actually supporting a system which causes worse moral outcomes is an empirical question? Defining the 'moral' outcomes is the philosophical one.

#### V

*EA owes its success as a philosophical-philanthropic movement largely to its eagerness and ability to work within existing political-economic systems.*

Ironically proving that it is possible to be successful with the system. (Sorry for getting a bit snarky to the end, but I honestly think this is just a bad article)

*Critical outlooks in which these ideas are at home have played no discernible role in discussions of EA, where there is rarely any suggestion of a tie between the forms of misery we are enjoined to alleviate and the structures of global capitalism. What is foregrounded instead is a paternalistic narrative about how the relatively wealthy should serve as benefactors of relatively poor and precarious humans and animals, and thus "do good."*

I guess because this isn't a true reflection on EA, and these views are likely to be considered and found to be wrong by most people who associate with EA?

*But it is also possible that the interventions of Effective Altruists will, because they affirm this system's underlying principles, contribute to its perpetuation, perhaps even precipitating the arrival of a further, more horrific "agricultural revolution." What is certain is that Effective Altruists' theoretical commitments lead them to approach animal protectionism without proper reference to political and economic forces that sustain factory farms. Anyone seeking substantial steps toward shutting down these "farms" would be well advised to exchange EA for efforts informed by an understanding of these forces. Only such interventions have a shot at being more than accidentally effective.*

It's possible, but how likely? Once again, these claims are both fundamental, all-encompassing, not-argued for, and not-supported.

*For EA to accommodate their voices, it would have to allow that their moral and political beliefs are in conflict with its guiding principles, and that these principles themselves need to be given up. To allow for this would be to reject EA in its current form as fatally flawed—finally a step toward doing a bit of good.*

I'm happy for EA to say that it isn't the right fit for everyone, and that you obviously don't have to be part of the EA movement to actually do good. Crary's last line is I think somewhat true in the first half, and the second half is just an unnecessary quip to end a poor piece.

— — —

## [17.The Change We Need](#)

*Lori Gruen*

*Comments:*

## [Coda—Future-Oriented Effective Altruism: What's Wrong with Longtermism?](#)

Comments: