

Mutual Aid as Effective Altruism

Ricky Mouser

Please cite the final version in the *Kennedy Institute of Ethics Journal*, vol. 33 no. 2, 2023, p. 201-226.

ABSTRACT: Effective altruism has a strategy problem. Overreliance on a strategy of donating to the most effective charities keeps us on *the firefighter's treadmill*, continually pursuing the next highest quantifiable marginal gain. But on its own, this is politically short-sighted. Without any long-term framework within which these individual rescues fit together to bring about the greatest overall impact, we are almost certainly leaving a lot of value on the table. So effective altruists' preferred means undercut their professed aims. Alongside the charity framework, the *more* effective altruist ought to consider a mutual aid framework, which better acknowledges and honors the unavoidably political commitments of effective altruism to reimagine and remake the world.

In recent years, effective altruists like Peter Singer and William MacAskill have recommended *earning to give* as a strategy for doing more good in the world.¹ Instead of becoming a charity worker, Singer argues, you should try to work in finance if you can:

There is usually no shortage of applicants for jobs with charities, so the charity will appoint someone else who will probably do almost as good a job as you would have done. "Almost" because if you had been offered the job, we can assume the charity considered you the best applicant for the position; but the difference between you and the next-best applicant is unlikely to be great. As a charity worker, therefore, you are largely replaceable. Working in finance, however, you earn much more than you need and give half of your earnings to the charity, which can use that money to employ two extra workers it would not, without your donation, have been able to employ at all. (Singer 2015, 41; compare MacAskill 2016)

¹ The New York Times reports that it was MacAskill who pulled Sam Bankman-Fried into effective altruism as an undergraduate. (Szalai 2022)

Obviously, a crisscrossing network of background assumptions underlies an argument like this. But in what follows, I'll argue that overreliance on the strategy of charitable donation is too politically short-sighted to be the most effective path forward in the long term.

Allow me to set the stage. In Section 1, I clarify the core commitments of effective altruism. In Section 2, I introduce a concern for effective altruists: Continually pursuing the next highest quantifiable marginal gain is too epistemically short-sighted to be the most effective long-term strategy. In Section 3, I distinguish mutual aid from the charity framework and highlight three key advantages of the former for political organizing. In Section 4, I tease out the underrecognized political commitments of *longtermism*, an important development within contemporary effective altruist circles. Finally, in Section 5, I argue that mutual aid can help achieve the long-term aims of effective altruism better than the charity framework alone by getting political enough to transform existing institutions.

To be clear: I'm not saying that mutual aid is the *most* effective form of altruism, or should be our *exclusive* strategy going forward. (I'm not sure how to argue that mutual aid alone beats out every other option on the table!) But I will make the case that in the long term, at least *supplementing* the charity framework with mutual aid efforts will be more effective at doing good in the world.

1. WHAT EFFECTIVE ALTRUISM IS

It's a matter of some controversy what the core commitments of effective altruism are. (e.g. Skelton 2016, Berkey 2018, MacAskill 2019) Singer cites a definition of effective altruism as “a philosophy and social movement which applies evidence and reason to working out the most effective ways to improve the world.” (Singer 2015, 4-5) But without clarification, this is far too thin to stake out a distinctive position. What counts as *applying evidence and reason*, or *improving the world*? So Brian Berkey helpfully distills the core commitments of effective altruism into four principles:

Strong Interest-Based Reasons: the needs and interests of individuals provide us with strong reasons to act so as to satisfy those needs and advance morally important interests. (Berkey 2021, 96)

That is, the interests of others—including animals—give us strong reasons to help them.

Importantly, these are other-interested reasons to act for *their* sake, not just self-interested reasons to act because (for instance) severe global inequality might compromise my own long-term financial stability.²

Cosmopolitan Impartiality: a global orientation, involving a commitment to consider the needs and interests of everyone as equally morally important, regardless of where they happen to live. (Berkey 2021, 96)

In “Famine, Affluence, and Morality,” Singer persuasively argues that distance is not a morally relevant factor so long as I can help someone ten yards or ten thousand miles away just as effectively. (Singer 1972, 231-232) Many effective altruists also think that *when* others happen to live is not a morally relevant factor so long as I can help someone alive now or in ten million years just as well. More on these so-called *longtermists* later.

Evidence-Based Decision Making: we should gather as much evidence as we (cost-effectively) can regarding what can be achieved through various efforts to improve the world, and, to the extent that we can, decide what to do in our own efforts to improve the world on the basis of that evidence. (Berkey 2021, 97)

In a previous formulation, Berkey emphasized the importance of employing “the best empirical research methods available.” (Berkey 2018, 147) But I think we should read the word ‘empirical’ broadly here, given that in the same paper Berkey acknowledges that effective altruists have leaned too hard on “quantifiable evidence of the kind that RCTs [randomized control trials] can provide.” (Berkey 2018, 160) Empirical evidence can come in many forms.

² Singer himself includes a curious appeal to the self-interest of the well-off: “In affluent nations a sizable segment of the population lives very comfortably and does not have to worry about economic security. In these circumstances, the need to find meaning and fulfillment in life comes to the fore, and many people turn to effective altruism as a way of giving their lives a purpose it would not otherwise have.” (Singer 2015, 94)

Unconditional Obligations of Beneficence: (at least) the well-off have at least some unconditional obligations to contribute to efforts to improve the world (Berkey 2021, 104)

Berkey notes that effective altruism does not merely make a conditional claim *on altruists* to give more effectively; it enjoins all those with (at least) substantial means to improve the world to do so, regardless of whether or not they would have recognized, upon reflection, some prior desire within themselves to be effective altruists. Further, these unconditional obligations are not trivial. Berkey argues that the well-off should give to at least the point of “non-negligible sacrifice.” (Berkey 2021, 104)

Berkey’s framework helps clarify the assumptions underlying Singer’s definition. Per Strong Interest-Based Reasons and Cosmopolitan Impartiality, *improving the world* involves acting upon our strong reasons to impartially consider and advance the needs and interests of all.³ And per Evidence-Based Decision Making, *applying evidence and reason* requires cost-effectively gathering and employing evidence to drive our interventions. Put together, effective altruism amounts to a project of impartially caring for all as best as we can reasonably know how.⁴

Given this framework, we can make better sense of Singer’s argument. As part of considering and working to advance the needs and interests of all, I ought to research various career paths for myself. When I do, I’ll notice that finance pays much better than charity. And, as Singer helpfully notes, “the more you earn, the more you can donate.” (Singer 2015, 39) So by earning to

³ Per *Unconditional Obligations of Beneficence*, these reasons are strong enough to render nontrivial self-sacrifice obligatory for at least some of us.

⁴ It might seem that Berkey’s revised framework, in aiming to be ecumenical, has lost track of some of the distinctive commitments of effective altruism, namely *welfarism* (welfare is the only ultimate moral value) and *strong maximization* (you ought to do the *most* good you can). I don’t think welfarism is a load-bearing pillar of the view, even if Oxford moral philosophers are often welfarists; making the good more capacious simply complicates the question of what we ought to do. And not all effective altruists hold that strongly maximizing (or even trying to) is a necessary condition for joining their ranks. Effectiveness as an altruist comes in degrees. So Singer distinguishes between living a *minimally acceptable* life, which “involves using a substantial part of our spare resources to make the world a better place,” and a *fully ethical* life, which “involves doing the most good we can.” He then acknowledges that “Most effective altruists are not saints but ordinary people like you and me, so very few effective altruists claim to live a fully ethical life. Most of them are somewhere on the continuum between a minimally acceptable ethical life and a fully ethical life.” (Singer 2015, vii, viii).

give, and giving to the most efficient charities, I can have a much greater impact than I would by working at a charity myself. The notion of *impact* is slightly technical, as Iason Gabriel explains:

A central claim of effective altruism is that for a person’s action to have impact, it must also be the case that her contribution would not have been made by someone else if she had chosen to act differently. (Gabriel 2017, 11)

The reasoning behind this marginal analysis framework is best described by MacAskill, who distinguishes between making a mere *change* (which someone else might have also made) and making a genuine *difference*. (MacAskill 2022, 42) What matters are the states of affairs that I bring about that wouldn’t have happened otherwise—the differences I make. If the same state of affairs would have obtained anyway, I’ve merely changed how it came to be without really making things better.

So if I become a charity worker, I could make a change, but not much difference; someone else would have done about as much good in that role. But if I take a job in finance, that’s a role that probably *wouldn’t* have been filled by someone giving half of their income away to highly effective charities. By snagging a job on Wall Street and donating my Christmas bonus, I can make a real difference. But charity work is just too fungible to be high-impact.

By working in finance, it’s true that I can send a lot of money to some impressively efficient charities. But where did my huge Christmas bonus come from? And is trying to become a banker just to give it all away really the life we want to recommend to the most morally conscientious people? While some effective altruist projects like 80,000 Hours have shifted towards recommending more morally neutral or beneficial jobs—though Srinivasan (2015) notes these include “quantitative hedge-fund trading, management consulting, [and] technology startups”—these larger, structural questions have still not received as much sustained scrutiny as they deserve from effective altruists themselves, even as their numbers have grown. There has been insufficient consideration of the fact that, as Pete Mills puts it, “professional philanthropy needs exploitation in order to mitigate the effects of exploitation.” (Mills 2012, 7)

2. THE FIREFIGHTER'S TREADMILL

Berkey describes effective altruism as both “a set of core philosophical commitments, and a social movement,” but it is a set of philosophical commitments first. (Berkey 2021, 93) Effective altruism begins by identifying our ends (states of affairs in which the world is improved) while deferring to our best empirical evidence about which means will bring those ends about. But even if effective altruism is formally “cause-neutral,” its practitioners time and time again suppose that the best way to improve the world is by donating to highly efficient charities that can demonstrably maximize the impact of each dollar they receive. (Gabriel 2017, 12) Thus, the means they tend to adopt involve “*giving* to efforts that will promote the relevant values most effectively.” (Berkey 2018, 147, my emphasis) The question that effective altruist-aligned metacharities like GiveWell take themselves to be answering is which charities deserve our funds, not whether the model of charitable donation is itself the best means of improving the world; *where* to give, not *how*.

In principle, effective altruists are open to other means of doing good, but data-driven charities can readily demonstrate the effectiveness of their targeted interventions with empirical evidence. However, as Emily Clough argues, these demonstrations can be less straightforward than they initially appear:

For example, an RCT might determine whether a bed net distribution program lowered the incidence of malaria among its target population. But it would be less likely to capture whether the program unintentionally demobilized political pressures on the government to build a more effective malaria eradication program, one that would ultimately affect more people. RCTs thus potentially miss broader insights and side effects of a program beyond its target population. (Clough 2015)

We’re beginning to see that quite a bit hangs on the nature and limitations of empirical evidence.

Evidence-Based Decision-Making requires that effective altruists cost-effectively gather and employ evidence about the world to guide their efforts to improve it. But we live in a complex world where

our charities function in “an inescapably political environment.” (Clough 2015) As a result, we may often be measuring the *changes* charities make when we think we’re measuring their *impact*.⁵

To be clear, highly effective charities can do amazing work. For example, according to GiveWell, the marginal cost of saving one life by donating to Helen Keller International’s Vitamin A supplementation program is about \$3,500. (GiveWell 2022) That’s a tremendous exchange rate for saving lives that should make us think twice before remodeling our kitchen or acquiring Twitter. But it doesn’t follow that donating to charities like Helen Keller International should be our exclusive or even our preferred strategy for making the world a better place.

After all, the charity framework comes with significant costs that are also worth recognizing. One, as Angus Deaton notes, is that “the world’s poor have such a passive role in all of this happiness creation.” (Deaton 2015) Charity supplements existing political and economic institutions without fundamentally challenging or threatening them; indeed, charity has an important role to play in maintaining things as they are. As Dean Spade argues, on the charity framework, nonprofits have to pitch themselves to well-off donors, leaving the needy themselves (and *their* conception of what they need) as an afterthought. (Spade 2020, 23; see also Syme 2019, 110) The charity framework offers well-off donors “a tax shelter,” “a lifestyle accessory,” and a way to “look generous while upholding and legitimizing the systems that concentrate wealth.” (Spade 2020, 25, 27, 23) In these ways, the charity framework helps to maintain the very systems that necessitate its existence.

There’s a real danger here of not seeing the forest for the trees. Within the charity framework, earning to give is arguably the highest-impact intervention I can make. But if we’re continually faced with a world full of people who desperately need rescuing, we should begin to wonder how we got here, and how we can jump off of what I will call *the firefighter’s treadmill*: a trap of

⁵ Syme goes so far as to argue that given its hidden costs, “charity may not be worth doing at all from a purely altruistic perspective.” (2019, 93)

continually pursuing the next highest quantifiable marginal gain instead of taking more epistemically opaque steps towards structural transformation. I'm not the first to identify this issue, but by calling it a *treadmill*, I want to emphasize that we may not actually be doing the most good we can this way, as effective altruists intend; instead, we might be running very fast, more or less in place, driving visible short-term change instead of ultimate long-term impact.

Rescuing the next person may appear more cost-effective than taking on the systems that continually imperil them. But there are two distinct epistemic challenges here. Amia Srinivasan notes how difficult it is to measure both the *probability* and the *value* of deep, systematic change:

What's the expected marginal value of becoming an anti-capitalist revolutionary? To answer that you'd need to put a value and probability measure on achieving an unrecognisably different world – even, perhaps, on our becoming unrecognisably different sorts of people. It's hard enough to quantify the value of a philanthropic intervention: how would we go about quantifying the consequences of radically reorganising society? (Srinivasan 2015)

Even if you think the expected marginal value will be negative, how can you tell? To calculate an expected value, we need to quantify the value of various outcomes and assign a probability measure to them, but both of these present significant methodological challenges.

First, it's unclear how to assign a probability measure if I'm only analyzing the marginal effects of my own individual actions. Critic after critic has noted that if I hold everything else constant except for *my* contributions, political transformation becomes more or less impossible. (e.g. Mills 2012, Srinivasan 2015, Herzog 2016) As a result, over-individualized marginal analysis systematically counsels against pursuing radical transformation and for working within the constraints of existing political and economic systems. At best, I can tinker within these systems, rounding off the sharp corners of the worst symptoms without ever addressing their deeper causes.

Despite some back and forth in the literature, (e.g. Berkey 2018, Dietz 2019, Berkey 2019) effective altruists themselves have begun addressing the question of what their movement hopes to accomplish together. In a keynote speech titled "Doing Good Together," William MacAskill

acknowledges that “Now, because there’s thousands of [effective altruists] around the world, we have the potential to do much more good acting in unison as a community than we can if we’re just acting as the sum total of a number of individual actors.” (MacAskill, 2017)

Coordination problems are, of course, inevitable. McMahan (2016) argues that in the end, we have to act as individuals, given whatever the facts are about others’ cooperation. But there’s not such deep disagreement on this point as it might appear. Whether we prefer an individual analysis that considers whether others will cooperate with *me*, or a collective analysis that asks how *we* should cooperate, there’s not a tremendous amount of functional difference between these competing views. The discussion about whether the best model of collaboration is individual or collective in nature seems to me a largely technical dispute.⁶

Either way, once we adequately account for coordination, the probability of transformative change will no longer hover quite so close to 0, so the expected value will rise. Suddenly, more interdependent efforts might attain plausibility. So, it’s only by starting to think about how we can work together to take on existing systems that we can get off the firefighter’s treadmill.

But even if we grant that an unrecognizably different world might be possible, it’s still unclear how to assign a value for *how much better* it might be, given the difficulty of measuring flourishing. It’s clear enough what extreme languishing looks like in terms of death, illness, starvation, illiteracy, and other absolute deprivations, and we can tally up these outcomes, more or less. But in a seriously unjust world where merely surviving is such a struggle for so many of us, we don’t have as clear a conception of what thriving looks like, much less how to quantify it.

⁶ I do want to highlight a small but important distinction between these two views: The collective analysis treats the whole (the moral community) rather than the part (the distinct individual) as the proper object of moral theorizing. Compare Srinivasan 2015: “[Effective altruism’s] utilitarian calculations presuppose that everyone else will continue to conduct business as usual; the world is a given, in which one can make careful, piecemeal interventions. The tacit assumption is that the individual, not the community, class or state, is the proper object of moral theorising.” Although I can’t pursue this further here, I suspect the collective analysis better captures the extension of morality beyond the self.

As a result, it's often easier to quantify the prevention of languishing than it is to quantify the promotion of flourishing. For instance, Psychologist Ed Diener's tremendously influential research program on subjective well-being outsources these judgments to individuals themselves. Consider his *Flourishing Scale*, created with Robert Biswas-Diener.⁷ Participants are presented with eight Likert scales and asked to give a response from 7 (Strongly agree) to 1 (Strongly disagree):

I lead a purposeful and meaningful life
 My social relationships are supporting and rewarding
 I am engaged and interested in my daily activities
 I actively contribute to the happiness and well-being of others
 I am competent and capable in the activities that are important to me
 I am a good person and live a good life
 I am optimistic about my future
 People respect me (Diener et al. 2010, 154-155)

“A high score,” the assessment concludes, “represents a person with many psychological resources and strengths.” (Diener et al. 2010, 155)

To my knowledge, no effective altruist metacharity has tried to make serious use of any of Diener's psychological inventories. Even if they did, this approach would be importantly incomplete. By farming out these judgments to individuals themselves, we paper over important differences in how they conceive of, e.g., being a good person and living a good life, and treat them as the uninterrogable authorities on their own lives, incapable of error or self-deception. Flourishing is a much more slippery empirical target than languishing.

This means it's hard for us to get a handle on how much better things could possibly be. And because we can't adequately measure flourishing, it's languishing that tends to dominate our Evidence-Based Decision Making. Chip Ward calls this *the tyranny of the quantifiable*: “what can be measured almost always takes precedence over what cannot.” (Solnit 2014, 104-105) As practical reasoners, we tend to give extra psychological weight to factors that are quantifiable, and diminished

⁷ Interestingly, while this was initially published as the Psychological Well-Being Scale, a year later it was republished as the Flourishing Scale. Compare Diener et al. 2009 with Diener et al. 2010.

weight to less quantifiable values.⁸ Even after repeated warnings to be more ecumenical about what might count as ‘empirical’ evidence, this bias is difficult to avoid. While everyone agrees that tools like quality-adjusted life years (QALYs) measure health imperfectly (to say the least) they do provide us numbers to go by. And so we often follow them, even to the point of fixation.

Within over-individualized marginal analysis frameworks, charitable donation has seemed the best path forward, and under the tyranny of the quantifiable, effective altruists have been very focused on improving measurable health outcomes in the short-term, though trends are shifting.⁹ But once we effectively pluralize our marginal analysis and recognize the tyranny of the quantifiable, we can see the firefighter’s treadmill for what it is: a short-sighted, piecemeal approach that jumps from one crisis to the next without any long-term framework for transforming (or adequately recognizing) the political and economic structures underlying them.

Of course, political and economic transformation is risky and opaque. It’s hard to know our chances of success, or how much better things might actually get, and assessing such risks inevitably draws upon controversial commitments about the nature and mechanisms of social change. Berkey is right that “there are good reasons for EA, as a social movement, to avoid adopting controversial philosophical positions (beyond those that constitute its core commitments) *where it can.*” (Berkey 2019, 332, my emphasis) Neutrality helps effective altruism pitch a big tent.

But trying to avoid such commitments simply because they are controversial amounts to an implicit political commitment to the values underlying the status quo. This keeps us on the firefighter’s treadmill, unable to adequately assess whether addressing problems at the root might bring about better outcomes. As Syme notes, “All actions within a system either perpetuate the

⁸ See Nguyen’s related work on value capture (Nguyen 2020, Chapter 9; Nguyen 2021) for more on how simplified, typically quantified metrics of our values can become so salient that they “take over” our practical reasoning, to our own detriment.

⁹ More recently, effective altruists have been less outspoken about earning to give as a central strategy of the movement, and as we’ll see, they’ve also become enamored with preventing x-risks in the long term.

status quo or seek to change it, so people have to decide which analysis they believe before they can identify a range of possible interventions to then experiment with and evaluate empirically.” (2019, 102) We have to take a stand on what we think the mechanisms of social change and the causes of current inequalities are, and these commitments can’t be outsourced to purely scientific inquiry.

Worse, downplaying our analysis of social dynamics comes with significant costs. Rescuing now may make us *feel* like we are making the most progress we can, but without a long-term strategy we’re just pursuing the next marginal gain wherever we can quantify it. Without any long-term framework within which these individual rescues fit together to bring about the greatest overall impact, we can end up leaving a lot of value on the table. A strategy of donating to the most effective charities I can find each year is not up to the task of fostering long-term political transformation. There must be a better way.

3. BEYOND A POLITICS OF PITY

It’s worth lingering over the language of rescue here, which traces back to Singer’s original discussion in which he imagines seeing a child drowning in a shallow pond. He concludes that he ought to rescue the child, even if it will mean muddying his clothes. (Singer 1972, 231) But since then, the heroic rhetoric has really come into its own. Jennifer C. Rubenstein notes that both Singer’s and MacAskill’s 2015 books specifically ask us to imagine *kicking in the door of a burning building* to save the lives of those inside, before triumphantly suggesting that we can do just as much good, or even more, by becoming effective altruists. (Rubenstein 2016, 520)

This language isn’t incidental, Rubenstein thinks; instead, she argues that it is the first lesson of effective altruism’s “hidden curriculum” that “teaches people in wealthy countries to think of themselves as heroes and rescuers.” (Rubenstein 2016, 520) There is a distinct moral danger here: I think this cheery saviorism papers over pity, as opposed to compassion.

These are two very different perspectives from which we might recognize our Strong Interest-Based Reasons. In *pity*, we look down on the other and feel sorry for them; in *compassion*, we put ourselves at another's level and suffer with them (com-passion). Where pity upholds my relative privilege, compassion foregrounds my horizontal concern for the other *as an equal*.

Effective altruism should nudge us to prefer the perspective of compassion over pity, given its core commitment to Cosmopolitan Impartiality, which again, requires us “to consider the needs and interests of everyone as equally morally important, regardless of where they happen to live.” (Berkey 2021, 96) But I think it's difficult to avoid concluding that the charity model reinscribes pity over compassion. It's not just (as we've already seen) that a select few get to finance and direct charity operations; the charity model fundamentally operates as an exogenous intervention, a top-down approach that identifies the suffering of others and works to plot and enact solutions *for* them.

But does this matter, given that the core concern of effective altruism is to do the most good? Here I think we need to introduce more *kinds* of goods to the conversation. One useful distinction is between material and existential goods. (See Mouser, forthcoming) The goal is not simply to maximize the good in material terms—e.g., the provision of food and medical treatment—even if these goods are the most readily quantified. We must also attend to *existential* or fundamentally respect-based goods as well.

But hierarchical models of aid administration reinscribe the logic of pity, which fails to fully respect the needy as our equals rather than mere patients whose outcomes are to be measurably improved. A more horizontal, *cooperative* model of aid would be more alive to the egalitarian respect that motivates the effective altruist commitment to Cosmopolitan Impartiality. And this matters, because the *quality* of aid with respect to how it functions and what it communicates to those aided about their status in the exchange, is itself a morally significant aspect of the good that effective altruists seek to do more of.

But recognizing that I am not a hero, merely a concerned equal, is already a distinctly political shift. By discarding the ethic of heroism, I insist that my role is to help empower and care for others, not to decide how to save them. This commitment to egalitarianism—to treating other persons as my equals—is best embodied by practices that continually challenge and flatten hierarchies with a logic of collaborative compassion rather than administrative pity. We need to attend to the value of working *with* those we aim to assist.

Arguably the *most* horizontal logic on offer is that of mutual aid, which Dean Spade analyzes as “collective coordination to meet each other’s needs.” (Spade 2020, 7) It’s important to emphasize that despite its reputation, mutual aid is not a fringe leftist strategy or anarchist sideshow; this notion includes a tremendously wide swath of practices. As Spade says:

We see examples of mutual aid in *every single social movement*, whether it’s people raising money for workers on strike, setting up a ride-sharing system during the Montgomery Bus Boycott, putting drinking water in the desert for migrants crossing the border, training each other in emergency medicine because ambulance response time in poor neighborhoods is too slow, raising money to pay for abortions for those who can’t afford them, or coordinating letter-writing to prisoners. These are mutual aid projects. (Spade 2020, 7, my emphasis)

The ubiquity and horizontality of mutual aid make it a natural place to turn, though it might almost seem *too* ubiquitous. But “collective coordination to meet each other’s needs” is not quite capacious enough to include the charity framework. When we listen and respond to others via mutual aid, we collaborate with the needy as our equal partners, not as helpless victims in need of rescue. By operating from the logic of compassion instead of pity, we acknowledge that *we* are fundamentally needy beings as well.

Let me head off an immediate objection: Mutual aid may seem to violate Cosmopolitan Impartiality by directing my immediate concern especially towards those within the community we are building together. There are two things to say here. First, I don’t think this involves seeing the needs or interests of those outside my community as less *important*. Directing my moral concern

especially towards my cousins does not require seeing their needs and interests as more important than the needs and interests of yours. Treating others *as equals* (moral peers) does not require treating them *equally* (identically).

Second, engaging in mutual aid may well be *part* of doing the most good overall. This is particularly true if it is more likely to lead to transformative political change in the long run, so that the needs and interests of everyone will be advanced more equitably in the future. Ultimately, *some* form of effective political organization is necessary to hop off the firefighter's treadmill, but the charity framework seems ill-suited to achieve this on its own. As a matter of strategic diversification, some measure of coalition-building at the community level is plausibly part of a more effective long-term strategy.

To motivate that thought, let's consider three key advantages of mutual aid over charity with respect to political organizing. By insisting upon horizontal structures of cooperation rather than vertical structures of charitable donation, mutual aid (1) mobilizes people, (2) expands solidarity, and (3) builds movements. (Spade 2020, 12)

First, mutual aid *mobilizes people* by politically empowering those who are aided without simply intervening on their behalf. Here our broader understanding of what counts as empirical evidence is directly relevant. The underlying commitments of mutual aid that "those on the front lines of a crisis have the best wisdom to solve the problems" (as opposed to waiting for impersonal standards of quantifiable evidence to be met) and that "collective action is the way forward" (rather than top-down intervention) are not borne from the evidence of randomized controlled trials but from the lived, first-personal experiences of those who have languished under seriously isolating and undemocratic institutions and noticed that their needs are not only ignored but actively thwarted by them. (Spade 2020, 7)

This too is empirical evidence in the sense that it is borne from the experiences of those who we have reason to see as domain experts; the oppressed themselves. (See Pearlman and Williams 2022) If we insist that only quantifiable or scientific evidence should count as empirical, we again run from the fact that the value commitments we have about which outcomes are themselves good, or even what counts as empirical evidence, are not at bottom scientific questions. These are questions about *values*. Trying to remain ‘neutral’ on such questions leaves the background values that structure and underlie the status quo unchecked.

This leads directly into the second point: Mutual aid *expands solidarity* by promoting democratic consensus at precisely the points where charity elevates technocratic expertise. Spade notes that on the charity framework, credentialed experts insist that they know best:

The idea promoted by nonprofits and universities is that people with advanced degrees are best suited to figure out the solutions to social problems. It mystifies the causes of poverty, making it seem like some kind of mysterious math problem that only people with advanced degrees can figure out. But any poor person knows that poverty is caused by the greed of their bosses, landlords, and health insurance companies, by systems of white supremacy and colonialism, and by wars and forced migrations. (Spade 2020, 26)

But in their rush to maximize quantifiable impact, effective altruists have failed to adequately analyze the causes of poverty, or to seriously consider whether that deprivation is structurally necessitated by our deeply undemocratic political and economic institutions, often dismissing such analyses too easily as *vague* or *motivated* (see Syme 2019). As a result, they stay trapped on the firefighter’s treadmill, and top-down strategies like earning to give seem doomed to keep them there by perpetuating these unjust systems.

On the other hand, mutual aid expands solidarity—our commitment to stand with one another, and deliberate and act together—precisely where the charity framework creates distinctions between donors, charity workers, eligible recipients, and ineligible recipients. This logic of vertical

administration preserves existing decision-making hierarchies about how society will be organized rather than fostering horizontal consensus and solidarity.

To see this, consider a charity like GiveDirectly, which extends the traditional mutual aid tactic of unconditional cash transfers across international lines. GiveDirectly takes on the mutual aid insight that we should defer to the wisdom of those on the ground to solve their own problems most effectively. We should treat *them* as the relevant domain experts. Even so, because all interaction between donors and beneficiaries occurs through the anonymizing intermediary of the charity framework, we miss out on what Savannah Pearlman calls the *relationship goods* of community building, solidarity, and coordinated, fully collective care. (forthcoming in this issue)

GiveDirectly succeeds at empowering the global needy financially, and goes some way towards subverting the charity framework by not acting *for* them. But it still fails to truly *connect* them with donors, to build a social or political community which bridges and transcends anonymous electronic transactions. GiveDirectly is a very effective charity, with empowering impacts on communities that are not fully captured by the metrics of effective altruist metacharities. But because it still makes all the decisions about where and how direct cash transfers will be sent, it's not truly horizontal mutual aid. Deferring to the experts is not *bad*, and may produce quite a bit of good relatively quickly. But it will fail to produce the relationship goods that power political movements in the long term, and so this strategy fails to build solidarity.

That brings us to the third and final point: Mutual aid *builds movements* that allow us to reimagine and remake our global political and economic landscape together, rather than indefinitely pursuing concretely realizable goods within the one we have. As we've seen, positive political transformation is only feasible if we work together. Solidarity not only makes it possible for us to see the bigger picture; it also allows us to succeed together:

Solidarity across issues and populations is what makes movements big and powerful. Without that connection, we end up in disconnected groups, working in their issue

silos, undermining each other, competing for attention and funding, not backing each other up and not building power. (Spade 2020, 15)

As we work together in mutual aid to meet each other's needs, we increasingly discover the interconnectedness of our struggles. Homelessness is driven by a lack of mental health services, unchecked rent increases, unlivably low minimum wages, shockingly common medical bankruptcies, and so on. These struggles are not only connected; they also have impacts on issues beyond homelessness. Coming together to practice mutual aid builds community awareness of these structural interdependencies, which is crucial, because if we're ever going to win, we have to win together.

But the charity framework doesn't bring us together. The effective altruist who bankrolls multiple charity employees takes a moment of their time each year to write a check for the most effective charities they can find, but otherwise, they're not systematically involved in building a live political movement. There's no larger vision of the world's political or economic institutions that they're helping to shape or actualize; their vision of the world is just one with a few more QALYs in it. Indeed, Rubenstein argues that the siloed charity framework can even be antithetical to solidarity-building, because the charity I donated to last year might become less effective or be surpassed in effectiveness by another. (Rubenstein 2016, 520-521) Chasing cause-neutral marginal value renders my next dollar *mercenary*.

Mobilizing people, expanding solidarity, and building movements are three important advantages of mutual aid over the charity framework, especially in the long term. I want to show that this is true even by effective altruists' own lights. But first, we'll have to see how many effective altruists themselves evaluate long-term considerations.

4. PLL SHOW YOU LONG TERM

This section is addressed to all effective altruists, but especially committed longtermists. We share more in common than you might think.

Longtermists care just as much about helping someone who will live ten million years from now as someone alive today. MacAskill’s “argument in a nutshell” for longtermism goes like this:

Future people count. There could be a lot of them. We can make their lives go better.
(MacAskill 2022, 16)

Even if we halt population growth at ten billion and maintain that level for the next five hundred million years, MacAskill calculates that there will be quadrillions of future persons. (MacAskill 2022, 23-26) Their sheer quantity ensures that, on an impartial calculus, concern for their needs will dwarf ours. It’s just math, and it will never work out in our favor.

Although longtermism is controversial, it’s becoming an increasingly influential position within the effective altruist community, presumably because it extends the core commitment of Cosmopolitan Impartiality not only across space but over time. (Call it *Cosmic* Impartiality if you like.) As a result, lots of attention has been directed towards mitigating existential risks (x-risks for short) like permanent AI enslavement or asteroid impacts. Nick Bostrom defines *x-risks* as “global, terminal risks...where an adverse outcome would either annihilate Earth-originating intelligent life or permanently and drastically curtail its potential.” (Bostrom 2002, 2) Longtermists argue that x-risks could impact so many future individuals that even a miniscule probability of averting disaster is worthy of significant investment.

This is particularly interesting because book titles like *The Most Good You Can Do* and *Doing Good Better* emphasize the aim of doing good, not preventing bad. Effective altruists’ marginal analysis framework appears to treat these as formally equivalent, because, assuming people’s lives are generally good for them, positive welfare has more time to accrue if we can avoid suffering serious x-risks. To this end, 80,000 Hours lists “Risks from artificial intelligence” and “Catastrophic pandemics” as the two most pressing problems facing the world today. (80,000 Hours, 2022)

But Srinivasan notes that once we zoom out like this, effective altruism can seem to lose contact with what made it attractive in the first place: the idea that there is so much preventable suffering all around us that things could be substantially improved *right now*. (Srinivasan 2015)

My diagnosis is that we need to expand our longtermist vocabulary. Alongside x-risks, I want to introduce *x-opportunities*, which are global, transformative opportunities where a favorable outcome would dramatically promote our flourishing. Ironically, x-opportunities have been insufficiently appreciated by an effective altruist movement obsessed with finding neglected causes where our contributions will have the highest marginal impact, although MacAskill does recognize that there are two ways we can impact future generations:

First, we can affect humanity's duration: ensuring that we survive the next few centuries affects how many future generations there are. That is, we can help *ensure civilisation's survival*...Second, we can affect civilisation's average value, changing how well or badly life goes for future generations, potentially for as long as civilisation lasts. That is, we can *change trajectory*, trying to improve the quality of future people's lives over the life span of civilisation. (MacAskill 2022, 44)

X-risks may either imperil our survival altogether, or threaten grave harm to our trajectory.

Symmetrically, x-opportunities may either secure our long-term survival, or promise transformative benefit to our trajectory.

But how should we balance survival against trajectory? It's a familiar thought when we consider end-of-life cases that *merely* surviving is not particularly valuable; it's thriving that counts. This underlies the popular deprivation account of death, according to which death is bad for the one who dies because it deprives them of the goods they would have received if they had continued to live. Nagel argues that

the value of life and its contents does not attach to mere organic survival: almost everyone would be indifferent (other things equal) between immediate death and immediate coma followed by death twenty years later without reawakening. (Nagel 1970, 74)

And MacAskill suggests that human extinction would only be very bad if the lives of future people were “good enough.” (MacAskill 2022, 201)

Because flourishing is what makes our lives good enough, promoting x-opportunities shouldn’t be an afterthought. By preventing x-risks, we try to raise the floor of how bad things could be. And by promoting x-opportunities, we try to raise the ceiling of how good things could be. If we’re serious about increasing our impact, we need to attend to both.¹⁰

To the longtermists: Our current political and economic institutions promote x-risks like climate change and nuclear annihilation. They also prevent x-opportunities, like the emergence of deeply democratic political and economic institutions. That’s bad now, but it’s even worse when we consider how many future people there could be. Once we recognize the importance of both x-risks and x-opportunities, longtermism should make us *more* focused on politically transforming institutions, not less.

Yet in MacAskill’s brief discussion of the importance of political activism, he restricts his focus almost exclusively to the expected value of voting. (MacAskill 2022, 246-247) He does briefly mention that you can contribute to political campaigns in other ways, and recognizes that conversations with family and friends about our values can increase political participation—once again, as measured by two studies on voter turnout.¹¹ (Green and McClellan 2020, Schein et al. 2021)

¹⁰ Here effective altruists might object their focus on preventing languishing is not merely a contingent epistemic result, but that preventing extreme languishing is *morally more urgent* than promoting flourishing. But I think this case is more difficult for effective altruists to make than it first appears, given their focus on doing the most good. As we’ve seen from Nagel, surviving at a neutral level of welfare is not obviously good for you. And merely *decent* living, with a marginally positive balance of flourishing over languishing, isn’t that much better. It’s flourishing that makes our lives positively worth living, as opposed to neutral or even bad for us. So why would doing the most good require prioritizing prevention of the bad over promotion of the good? An asymmetry thesis here would require further argument and clarification. For example, concerns about the distribution of benefits and harms might be captured by considering existential goods, but assessing such a case lies beyond the scope of my efforts here.

¹¹ See also Pete Mills’s observation that the effective altruist project 80,000 Hours understands us to live “in a world of policy advisers, salaried campaigners, and party leaders, [where] politics is just another career.” (Mills 2012, 8)

But his abiding electoralism is striking, recalling Srinivasan’s chief complaint that effective altruists like MacAskill fail to challenge power and instead leave everything as it is:

MacAskill does not address the deep sources of global misery – international trade and finance, debt, nationalism, imperialism, racial and gender-based subordination, war, environmental degradation, corruption, exploitation of labour – or the forces that ensure its reproduction. Effective altruism doesn’t try to understand how power works, except to better align itself with it. In this sense it leaves everything just as it is. This is no doubt comforting to those who enjoy the status quo – and may in part account for the movement’s success. (Srinivasan 2015)

To avoid the firefighter’s treadmill, we have to get political. Failure to recognize this is oddly short-sighted from longtermists openly speculating about whether civilization might survive “beyond a million trillion years” if we can figure out how to harness the energy of black holes. (MacAskill 2022, 285) Even if the particular institutions we come up with now, like Monroe County Mutual Aid or Mutual Aid Disaster Relief, won’t last that long, these more democratic institutions allow us to better recognize and address everyone’s needs and interests, and give us a something to iterate on in the long term. In what follows, I’ll argue that mutual aid is better-suited to meet the needs and concerns of dedicated longtermists than the charity framework.

5. MUTUAL AID AS EFFECTIVE ALTRUISM

Assessing long-term impact is very difficult: How can I reliably predict how my actions will impact others indefinitely far forward in the future? Given this, MacAskill gives fellow longtermists three recommendations for what we should do right now. When facing such “empirical and evaluative complexity,” we should:

First, take actions that we can be comparatively confident are good...Second, try to increase the number of options open to us...Third, try to learn more. (MacAskill 2022, 239)

I admire these recommendations, and think mutual aid is well-positioned to satisfy them. So let’s take them in turn.

First, we should take actions that we can be comparatively confident are good. So sure—go ahead and take the time to cast an informed ballot. But we can confidently do better, and we must, because politics doesn't just happen at the ballot box. Spade highlights “the undemocratic infrastructure of our lives—the extractive and unjust energy, food, health care, and transportation systems.” (Spade 2020, 20) To do better in the long run, we have to survive now, and mutual aid has an important role to play as we figure out how to meet each other's present needs. Building communities of solidarity is a robustly good start, and a proof of concept where we can experiment at smaller scales and try out new ideas.

But how well will these ideas scale up beyond their communities of origin? Here I want to subtly shift ground from asking how mutual aid practices would look at global scale (Perhaps that *would* have to look more like GiveDirectly?) to asking how mutual aid communities can interact. After all, we currently lack overarching institutions that effectively *unite* us across the global community—the closest we have now is arguably the carefully curtailed United Nations. Under the charity framework, most of what we can point to are one-sided affairs in which one relatively privileged group more or less unilaterally decides upon an intervention to be visited upon another for its benefit. But mutual aid groups *do* interact effectively already: witness collaboration between The Movement for Black Lives and Pro-Palestinian Movement, (Wang 2021) or the stunningly wide coalition-building of Mutual Aid Disaster Relief. (Mutual Aid Disaster Relief 2023) Instead of searching for one-size-fits-all solutions that can scale up indefinitely, we can start by promoting horizontal dialogue and cooperation between smaller communities that are more sensitive to their own needs.

This leads us to MacAskill's second recommendation: Try to increase the number of options open to us. He recognizes that because our contemporary values are almost certainly not finished

products, we ought to cultivate diversity for the sake of moral progress. So he praises charter cities, and worries about premature value lock-in under one world culture or government:

a morally exploratory world would favour *political experimentalism*—increasing cultural and intellectual diversity, if possible...If we are aiming to get to the best possible society, we should worry about premature convergence, like a teenager marrying the first person they date. (MacAskill 2022, 108)

By insisting on horizontal decision-making and treating those on the ground (wherever they are) as experts, mutual aid encourages exactly the sort of diversity of thought and practice that MacAskill recommends. No two mutual aid groups operate exactly alike, nor should they, as Spade clarifies:

There is no one correct or perfect group culture. Groups should be different from each other because the people in them are different and we all bring different qualities, skills, and viewpoints. (Spade 2020, 71-72)

Effective altruists need not be committed to the view that diversity *itself* is an unalloyed good, given that at least some values and practices are morally superior to others. But MacAskill emphasizes the instrumental value of diversity for his third recommendation: learning more. By trying out diverse approaches in different communities, we can reconsider what our normative values ought to be *and* discover what works instrumentally. That is, there are two distinct kinds of x-opportunities here. We can clarify and improve our value commitments, and we can test and iterate practices that try to enact them by including more and more people.

Note that mutual aid, with its horizontal sensitivity to diverse views and practices, is better positioned to promote both kinds of x-opportunities than the charity framework. That's because building communities of solidarity is a superior vehicle for exploring the space of values, exploring the space of practices, and ultimately achieving political transformation. Instead of jumping ahead to administrating solutions, the mutual aid framework lets *everyone's* voice be heard, and as already noted, directly teaches us about the intersectionality of our struggles.

As Spade notes, mutual aid both “exposes the failures of the current system and shows an alternative.” (Spade 2020, 13) It gives us lived experience of what real democracy (people-power)

looks like: We tell you what we need, you tell us what you need, and we figure out together how to care for everyone as best as we can reasonably know how. Mutual aid allows us to reimagine what our local and global political and economic landscape could be like, and starts making it real before our eyes. Perhaps most importantly, solidarity teaches us how those who are different from us are also the same:

By working together and participating in shared political education programs, members could learn about experiences different from theirs and build solidarity across those differences. This changed—and continues to change—not only the individuals in the group, but the kind of politics the group practices. (Spade 2020, 14)

We've seen the limitations of the charity model, which doesn't just leave the unjust institutions of the world as they are but can actively bolster them. The mutual aid framework does better here by effective altruists' own lights, by beginning to mobilize people, expand solidarity, and build movements. This allows us to do robust good now, and in the long term, it promotes x-opportunities by increasing our options and allowing us to learn from each other's experiences. It's a much more holistic approach to pursuing the political and economic transformation needed to create a better world than staying on the firefighter's treadmill.

What would this look like in action? Earlier, I mentioned that 80,000 Hours considers artificial intelligence to be the most pressing issue of our time. I'm convinced by enough of Nick Bostrom's book *Superintelligence* to think that artificial intelligence really will confront us with enormous x-risks (and x-opportunities) much sooner than most people think. But even granting the diagnosis, what's the best cure? I don't think we should endlessly slush money towards a small cadre of Silicon Valley tech leaders. Funded and confronted daily by the values of venture capitalism, these figures are too near to the pursuit of exponential quarterly growth, and too far from the values of the rest of our communities. And they are only partly balanced by rapidly growing effective altruist think tanks and research programs, which provide another stratum of technical experts.

If powerful artificial intelligence really will capture trillions of dollars of the world's wealth, as OpenAI founder Sam Altman casually speculates, he plans to ask the artificial intelligence for help redistributing the world's captured wealth. (Metz 2023) *That* is the logical technocratic conclusion of the charity framework. But shouldn't we all get a democratic say in articulating what our values are and what sort of world we want to build together? Even if the machines become far 'smarter' than us, the prospect of unreflectively *outsourcing* all our rich and varied values to an unelected elite of industry-leading technical experts, or their morally opaque authorities like a hypothetical "Google Morals," raises the tremendously alienating x-risk of overriding our practices of valuing themselves. (Howell 2014)

Instead, we need to build communities of solidarity now where we can meet our present needs and begin articulating and exploring our values together through a process of horizontal, *democratic* inquiry. We need to foster contextually-sensitive conceptions of the good that grapple with the fundamental depth, complexity, and incommensurability of the things our specific communities value. And we need to build enough democratic power to ensure that the benefits and opportunities of powerful artificial intelligence are equitably distributed. We might not need to *abandon* the charity framework, but we do need to at least supplement and challenge it with mutual aid.

The effective altruist framework is committed to a far more politically radical project than its practitioners seem to realize. To become *more* effective altruists, they would do well to take a long, hard look at the democratizing power of mutual aid.¹²

¹² Thanks to Eric Alt, Rich Cameron, Binh Doan, Paul Howatt, Savannah Pearlman, Elizabeth Williams, a wonderful audience at the Prindle Institute, and my anonymous reviewers for their helpful comments and discussion.

Works Cited

- 80,000 Hours (September 2022). What are the most pressing world problems?
80000hours.org/problem-profiles/.
- Berkey, Brian (2018). The Institutional Critique of Effective Altruism. *Utilitas* 30 (2):143-171.
- Berkey, Brian (2019). Collective Obligations and the Institutional Critique of Effective Altruism: A Reply to Alexander Dietz. *Utilitas* 31 (3):326-333.
- Berkey, Brian (2021). The Philosophical Core of Effective Altruism. *Journal of Social Philosophy* 52 (1):93-115.
- Bostrom, Nick (2002). Existential Risks. *Journal of Evolution and Technology* 9 (1):1-31.
- Bostrom, Nick (2014). *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press.
- Clough, Emily (2015). Effective Altruism's Political Blind Spot. *Boston Review*
www.bostonreview.net/articles/emily-clough-effective-altruism-ngos/.
- Deaton, Angus (2015). Response to Effective Altruism. *Boston Review*
www.bostonreview.net/forum_response/response-angus-deaton/.
- Diener, Ed; Wirtz, Derrick; Tov, William; Kim-Prieto, Chu; Choi, Dong-won; Oishi, Shigehiro; and Biswas-Diener, Robert. (2009). New Measures of Well-Being: Flourishing and Positive and Negative Feelings. *Social Indicators Research* 39, 247–266.
- Diener, Ed; Wirtz, Derrick; Tov, William; Kim-Prieto, Chu; Choi, Dong-won; Oishi, Shigehiro; and Biswas-Diener, Robert. (2010). New Well-Being Measures: Short Scales to Assess Flourishing and Positive and Negative Feelings. *Social Indicators Research* 97, 143–156.
- Dietz, Alexander. Effective Altruism and Collective Obligations. *Utilitas* 31 (1):106-115.
- Gabriel, Jason (2017). Effective Altruism and its Critics. *Journal of Applied Philosophy* 34 (4):457-473.

- GiveWell (accessed Dec 22, 2022). Our Top Charities. www.givewell.org/charities/top-charities.
- Green, Donald P. and McClellan, Oliver A. (2020). The Effects of Election Day Festivals and Early Voting Center Festivals on Voter Turnout: A Field Experiment Conducted During the 2018 Midterm Election. *SSRN* ssrn.com/abstract=3548996.
- Herzog, Lisa (2016). Can Effective Altruism Really Change the World? *OpenDemocracy.net* www.opendemocracy.net/transformation/lisa-herzog/can-effective-altruism-really-change-world.
- Howell, Robert (2014). Google Morals, Virtue, and the Asymmetry of Deference. *Noûs* 48 (3):389-415.
- MacAskill, William (2016). Banking: The Ethical Career Choice. In David Edmonds (ed), *Philosophers Take on the World*. Oxford: Oxford University Press, 84-86.
- MacAskill, William (2015). *Doing Good Better: Effective Altruism and a Radical New Way to Make a Difference*. London: Guardian Faber Books.
- MacAskill, William (2017). Doing Good Together. Centre for Effective Altruism, www.youtube.com/watch?v=DQbrtJYWPIU.
- MacAskill, William (2019). The Definition of Effective Altruism. In Hilary Greaves and Theron Pummer (eds), *Effective Altruism: Philosophical Issues*. Oxford: Oxford University Press, 10-28.
- MacAskill, William (2022). *What We Owe the Future*. New York: Hachette Book Group.
- McMahan, Jeff (2016). Philosophical Critiques of Effective Altruism. *The Philosophers' Magazine* 73:92-99.
- Metz, Cade (2023). The ChatGPT King Isn't Worried, but He Knows You Might Be. *New York Times* <https://www.nytimes.com/2023/03/31/technology/sam-altman-open-ai-chatgpt.html>.
- Mills, Pete (2012). The Ethical Careers Debate. *The Oxford Left Review* (7), 4-9.

- Mouser, Ricky (forthcoming). How to Read a Riot. *The Journal of Ethics and Social Philosophy*.
- Mutual Aid Disaster Relief (accessed 2023). Co Conspirators. mutualaiddisasterrelief.org/co-conspirators/.
- Nagel, Thomas (1970). Death. *Noûs* 4 (1):73-80.
- Nguyen, C. Thi (2020). *Games: Agency as Art*. New York: Oxford University Press.
- Nguyen, C. Thi (2021). How Twitter Gamifies Communication. In Jennifer Lackey (ed), *Applied Epistemology*. New York: Oxford University Press, 410-436.
- Pearlman, Savannah (forthcoming in this issue). Solidarity Over Charity: Mutual Aid as a Moral Alternative to Effective Altruism.
- Pearlman, Savannah and Williams, Elizabeth (2022). Why You Ought to Defer: Moral Deference and Marginalized Experience. *Feminist Philosophy Quarterly* 8 (2).
- Rubenstein, Jennifer C. (2016). The Lessons of Effective Altruism. *Ethics and International Affairs* 30 (4):511-526.
- Schein, Aaron; Vafa, Keyon; Sridhar, Dhanya; Veitch, Victor; Quinn, Jeffrey; Moffet, James; Blei, David M.; Green, Donald P. (2021). Assessing the Effects of Friend-to-Friend Texting on Turnout in the 2020 U.S. Presidential Election. *Proceedings of the Web Conference 2021*, 2025-2036.
- Skelton, Anthony (2016). The Ethical Principles of Effective Altruism. *Journal of Global Ethics* 12 \ (2):137-146.
- Szalai, Jennifer (2022). How Sam Bankman-Fried Put Effective Altruism on the Defensive. *New York Times* www.nytimes.com/2022/12/09/books/review/effective-altruism-sam-bankman-fried-crypto.html.
- Singer, Peter (1972). Famine, Affluence, and Morality. *Philosophy and Public Affairs* 1 (3):229-243.
- Singer, Peter (2015). *The Most Good You Can Do: How Effective Altruism Is Changing Ideas About Living*

Ethically. New Haven and London: Yale University Press.

Solnit, Rebecca (2014). *Men Explain Things to Me*, Chicago: Haymarket Books.

Spade, David (2020). *Mutual Aid: Building Solidarity During This Crisis (and the Next)*. London and New York: Verso.

Srinivasan, Amia (2015). Stop the Robot Apocalypse. *London Review of Books* www.lrb.co.uk/the-paper/v37/n18/amia-srinivasan/stop-the-robot-apocalypse.

Syme, Timothy (2019). Charity vs. Revolution: Effective Altruism and the Systemic Change Objection. *Ethical Theory and Moral Practice* 22 (1):93-120.

Wang, Hansi Lo (2021). The Complicated History Behind BLM's Solidarity With The Pro-Palestinian Movement. *National Public Radio* www.npr.org/2021/06/07/1003872848/the-complicated-history-behind-blms-solidarity-with-the-pro-palestinian-movement.