

Stanford SOCIAL INNOVATION^{Review}

Viewpoint
In Defense of Pet Causes
By Ian David Moss

Stanford Social Innovation Review
Spring 2017

Copyright © 2017 by Leland Stanford Jr. University
All Rights Reserved

VIEWPOINT

In Defense of Pet Causes

The effective altruism movement could be more effective if it encouraged adoption of its principles within causes and geographies, not just across them.

BY IAN DAVID MOSS

While all philanthropy seeks to make the world a better place, the effective altruism movement believes this is not enough. Its leaders are trying to raise an army of analytical thinkers who use evidence and reason to work out how to do the most overall good possible. By any reasonable measure, I should be a prime target for conscription. I have spent my entire career in the nonprofit sector, except for the two years I took earning my MBA. I thrive on abstract debates about moral philosophy and daydream about how I would design my own foundation. I am the kind of person who draws up a theory of change for how to live my own life—and actually finds that exercise useful.

So why have I not yet signed up to fight in the trenches with my geeky do-gooder comrades? Well, for one, I have at times felt that the movement is more invested in scolding me for my lack of commitment to being an absolute saint than affirming the charitable impulses I do have. To cite just one example, effective altruist leaders such as philosophers William MacAskill and Peter Singer have argued that choosing lower-impact charities over ones that directly save lives is morally equivalent to killing people.

Even if you agree with that argument, your position is likely to make you an outlier. A wide swath of research on donor motivation indicates that giving typically has much more to do with factors such as having a connection to a specific cause, feeling an obligation to “give back,” and cultivating social relationships rather than dispassionate explorations of the best opportunities to make a difference.

According to a landmark 2010 donor segmentation study by Hope Consulting that focused on affluent individuals in the United States, only 4 percent of donors consider the effectiveness of an organization the key driver of a gift, and just 3 percent actually research organizations’ effectiveness in order to choose which one to support. The William and Flora Hewlett Foundation not long ago decided to shut down a program intended to increase donors’ demand for information in large part because such study results, along with a subsequent evaluation of its efforts in this arena, were so discouraging.

Effective altruism’s growth in recent years has been impressive, but even if donations to charities approved by GiveWell or other effective altruist organizations were to reach as much as \$1 billion annually, that figure

would represent less than one-third of 1 percent of total donations in the United States alone. According to MacAskill, effective altruists “want to make the *most difference*” rather than merely “*a difference*” in the world. Taking him at his word, then, we should give serious credence to the idea that it would be more effective, all things considered, for the movement to encourage the adoption of effective altruist ideas *within domains*.

Domain-specific effective altruism would urge that principles of effective altruism be followed within an area of philanthropic focus, such as a specific cause or geography. So if you are especially concerned about, say, Seattle, you would apply effective altruist principles to try to figure out the absolute best and most cost-effective ways to make Seattle a better place. Similarly, if what you care about most is climate change, domain-specific effective altruism would focus on helping the climate change field make more meaningful progress faster.

BINDING COMMITMENTS

Consider my wife and me. Several years ago, not long after we got married and opened our first joint bank account, we sat down together to plan our giving budget. I came to the conversation ready to advocate for

a large portion of our philanthropy to go toward providing anti-malaria bed nets for impoverished people in Africa. My interest in the Against Malaria Foundation derived not from any particular expertise of mine in public health or experience living in malaria-ridden countries, but rather from my fascination with GiveWell, an organization respected by effective altruists for its rigorous, evidence-driven charity recommendations.

But my wife took a different perspective. For her, charity begins locally—with the people whose lives



IAN DAVID MOSS (@iandavidmoss) is vice president of strategy and analytics at Fractured Atlas and founder of Createquity, a think tank and online publication investigating the most important issues in the arts and what we can do about them.

intersect ours every day. To do nothing to acknowledge our privilege and our neighbors' lack of it would be callous and dehumanizing. It was crucial for her that, whatever other decisions we made, we would reserve some of our money for those in our city who need it.

My wife was not the only one who had strong preferences for our giving. My professional life has been spent almost entirely in arts and culture—hardly a priority for effective altruists. For the past decade, I have operated a website devoted to determining the most important issues in the arts and what we can do about them. To abandon the arts in my charitable budget would have felt like a denial of a core element of my identity. If people as passionate about the arts as me refused to make them a giving priority, who would step in?

In the end, we carved out about half of our charitable giving budget for these special priorities, which neither of us really saw as negotiable. But that still left open the question of which specific organizations we should support in our preferred domains. Unlike with our gift to fight malaria, there was no GiveWell for arts charities or local social service organizations to guide us in that decision. So we simply did the best we could with the limited time we had, undoubtedly leaving opportunities for impact on the table.

I suspect that my wife and I are not alone. Most criticism of effective altruism comes from people who reject its demand to stray from favored causes and geographies. By contrast, I see far less contention with effective altruism's core mandate to try to do the most good possible with one's contributions. On the contrary, if the conversation is about a topic in which they are invested, many donors are *very* interested in learning about the most effective ways they can make change.

A DIVERSIFIED PORTFOLIO

Strict effective altruists would object to these types of domain-specific restrictions as less virtuous than operating with none at all, and I do not disagree. But all of us have limits to our charity. No one, Peter Singer included,

forgoes every self-indulgence just because the resources involved could help someone else in need. The error of effective altruists is to see all self-indulgences, even charitable donations within lower-priority cause areas, as entirely outside the scope of what matters to effective altruism. Discounting the relevance of these passions risks making effective altruism itself irrelevant to most givers.

Domain-specific effective altruism, by contrast, could more effectively engage more donors and do-gooders by persuading them to adopt some effective altruist principles without losing them from the start by dismissing their favored causes. Surely if the strict effective altruist is forced to choose, she would prefer that people make the greatest possible impact within domains they care about (but that may not have the highest potential for overall impact) over not only sticking to lower-potential domains but failing to make a substantive difference in them.

Domain-specific effective altruists need not reject the definition of “the most good” used by effective altruists who take a fully global perspective. Indeed, coordinating goals across focus areas is what distinguishes domain-specific effective altruism from ordinary strategic philanthropy. For example, Createquity, the organization I founded that researches the most important issues in the arts, makes an explicit connection between the arts and a broader conception of overall well-being—the same thing that many effective altruists are trying to maximize through their efforts. In fact, one of Createquity's core principles is Net Benefit, which reads, in part: “We don't ever want to be in the position of supporting the arts at the expense of the rest of society.” Thus, any success Createquity has in identifying promising issues within the arts field and motivating productive action on them will be consistent with improving outcomes from a global effective altruist perspective.

Finally, embracing domain-specific effective altruism also has the advantage of diversifying the portfolio of potential impact for effective altruism. Even within the effective altruist movement, there are disagreements

about the highest-potential causes to champion. Take the domain-specific effective altruist organizations that, arguably, already exist, such as Animal Charity Evaluators (ACE) or the Machine Intelligence Research Institute. Animal welfare and the development of “friendly” artificial intelligence (AI) are both considered causes of interest for the effective altruist movement. But how should they be evaluated against each other? And more to the point, if it were conclusively determined that AI safety was the optimal cause to focus on, would ACE and other effective animal charities shut down to avoid diverting attention and resources away from AI safety?

The reality, as most effective altruists will admit, is that virtually all estimates of the expected impact of various interventions are rife with uncertainty. A Wall Street investor would be considered insane to bank her entire asset base with a single company or industry, and if anything, estimates and predictions are even less precise in the social realm than they are in business. Just because bed nets have been estimated to save lives in the past does not guarantee that your donation to an anti-malaria organization today will end up being the best or even a good use of your money. What if the evidence underlying movement dogma turns out to be flawed? What if there are better causes that have not been properly researched yet? Effective altruists are right to seek to prioritize the most promising opportunities, but it is far from clear that those opportunities all reside in causes currently favored by their community.

Effective altruists have a truly revolutionary idea, yet their own rhetoric and ideology are currently limiting that potential. The few heroes who are prepared to embrace any cause in the name of global empathy should be treasured and cultivated. But solely relying on them to change the world is unwise. If effective altruists fail to engage those who want to maximize their impact but will not abandon causes and geographies they care about deeply, their army of do-gooders will have a hard time winning the war for hearts and minds. ■