



## Is Your Charity “Effective?” Effective Altruism and Your Donors

PG Calc Featured Article, April 2023 eRate

By Kara Morin, PG Calc Director of Client Services

Recently, researchers at the Harvard Center for Brain Science crafted an [experiment](#) to see if donors could be redirected from supporting their favorite charities (“giving from the heart”) to supporting more effective charities (“giving from the head”). The effective charities were selected from [GiveWell](#), a charity navigator for the [Effective Altruism](#) (EA) movement, which also funded the study through the Effective Altruism Fund. If your only familiarity with EA is its relationship to disgraced crypto-currency billionaire Sam Bankman-Fried, it’s worth learning a bit about the movement and how it’s working to change the nature of philanthropy.

EA encourages its followers to *earn more, so they can give more* and to *support only the most effective charities* with their giving. According to EA, an effective charitable gift benefits the most people for the least financial outlay. For instance, let’s say a donor is interested in making a gift that will help the blind. They could make a gift to train a seeing eye dog for \$50,000 and benefit a single person. But if they were influenced by EA, they might decide instead to pay for two cataract surgeries at \$25,000 apiece. From a purely numbers-based perspective, the second gift is twice as effective.

But EA also asks donors to consider all the people who **might** become blind at some time in the future, and whether it might be **most cost effective** to make a gift that helps them avoid blindness. If antibiotics for trachoma (a disease which sometimes causes blindness) cost \$50 a prescription, EA would say it is more effective to minimize the chance that 1,000 people might become blind than to restore sight to two people. The probability of their blindness doesn’t factor into the equation: 1,000 is more than 2, so EA says choose the antibiotics.

This future focus is sometimes called “long-termism.” At an extreme it works something like this: mathematically speaking, there will be more people in the future than exist today, so if a donor truly wants the best ROI for their charitable dollar, they should focus their giving on existential threats to humanity. For example, gifts to get humans settled off of Earth before our sun dies 10 billion years from now would be effective gifts. If that sounds like an unfair characterization, rest assured that EA is already being used to make a moral case for space settlements on Mars.

While GiveWell continues to focus on high-volume, low-cost charitable gifts to fund the distribution of mosquito nets and vaccines, its EA partner [GivingWhatWeCan.org](https://givingwhatwecan.org) identifies research into artificial intelligence safety as one of the most effective ways of giving, since it will help humanity avoid the *Terminator* scenario.

So, if your charity doesn't distribute mosquito nets or conduct space exploration, do you need to worry that EA will capture your donors? They're certainly going to try.

The first phase of the Harvard experiment took a purely numbers-driven approach. Participants were given \$100 to donate to any 501(c)3 but were informed before making the gift that their gift would be 100 times more effective if given to a charity selected by GiveWell. Participants were largely unpersuaded by this EA approach; 83% used their \$100 to support their favorite charity.

In the second phase of the experiment, participants were encouraged to divide their gift between their preferred charity and an effective charity selected by GiveWell. If participants agreed to divide their gift, they received a match. If the division of the gift favored the effective charity (such as 60% effective/40% preferred), the greater the matching gift amount. The result was that *half of participants decided to divide their gift* between their favorite charity and an effective charity.

The Harvard paper is fascinating reading, not least because it illustrates that the road to changing how donors think about giving could be uphill. The Harvard researchers summarized the challenges in their abstract:

“First, many people are unaware of the massive differences in effectiveness across charities. However, even when informed of such differences, few donors are motivated to give based on effectiveness. Likewise, few are willing to pay to learn about the effectiveness of charity options. Providing effectiveness information can even reduce giving, and some donors may use effectiveness measures as an excuse not to give. Rather than focusing on effectiveness, most donors favor charities that are personally meaningful and emotionally appealing.”

But is there anything inherently wrong about gifts that are personally meaningful rather than numerically rational? Is it always better to prevent suffering in the future if it's at the expense of relieving suffering in the present? What makes a gift truly effective? On these

questions the Harvard study is inconclusive, “Whether it is good for donors to give more effectively in this sense is a value judgment that goes beyond the scope of this paper.”

Ultimately, the Harvard study is most useful as a window into what EA wants to do, which is change the nature of charitable giving. And, because redirecting existing donors itself is seen as an effectiveness multiplier, this type of research will continue.