

What you get for your dollar

“Effective Altruism” — What it is, how philanthropic foundations use it and what are its risks and side-effects

by Karolin Seitz

In parallel to the debate on ‘Aid Effectiveness’ among donor and recipient countries, a new approach, labeled “Effective Altruism” (EA) has gained traction in the debate on the impact of development aid. Unlike Aid Effectiveness, which involves donor and recipient governments, the new approach involves philanthropic institutions and their chosen (non-)governmental recipients and explores ways in which funding can be used most efficiently to have the greatest impact. Its proponents, including new philanthropic entities and so-called ‘venture philanthropies’ such as the Bill & Melinda Gates Foundation, claim that their funding decisions are based on evidence-based results. However, such decisions are primarily grounded in cost-benefit considerations, neglecting social and cultural considerations and looking at problems in isolation from the wider context.

Ensuring that charitable giving is used for the maximum good, is a worthy goal. However, from a human rights perspective, several concerns arise concerning the underlying assumptions, the methodology and the consequences of the practical application of EA.

This briefing paper provides an overview of the approach underlying EA, how and by whom it is applied and its problems and consequences. It concludes that policy makers, rather than be guided by its assumptions and conclusions, must instead concentrate on understanding the confounding structural causes of interdependent global challenges and aim at their long-term solution, within an overarching human rights framework.

Since the beginning of international development assistance there has been intense debate throughout the international donor community as to if, how and which forms of development aid can actually have the greatest and most sustainable impact on people’s lives. In the last two decades, governments from donor and recipient countries met four times to discuss these issues, summarized under the theme of “Aid Effectiveness”. Meeting in Rome (2003), Paris (2005), Accra (2008) and Busan (2011), governments committed to a set of principles and joint actions that aim to enhance effectiveness of development cooperation. In 2012, governments, civil society organizations and actors from the private sector established the Global Partnership for Effective Development Cooperation, a multi-stakeholder forum to strengthen their respective efforts.

Since then, however, this approach to Aid Effectiveness has lost political momentum.

In recent years, a new approach, labeled “Effective Altruism” (EA) has asked for attention in the debate on the impact of development aid. It is no longer focused on donor and recipient governments, but instead on individual donors, philanthropic institutions and their chosen charitable organizations. Players seek to determine how donations could be used most efficiently—in terms of costs and time—to have the greatest impact.¹

¹ The label “Effective Altruism” is somewhat misleading. In fact, the proper label would be “Efficient Altruism”, as it clearly is centered around how to achieve the greatest good with the smallest input possible. Were it to concentrate on effectiveness, it would concentrate not on input, but exclusively on outcomes.

Proponents claim that their funding decisions are based on empirical evidence and rational thought rather than on allegedly ‘emotional’ decision making.² Under the premises of “Effective Altruism”, impact does not depend on what benefit is being achieved, but what difference it makes in the overall funding landscape. If the individual does not donate for a specific purpose, would someone else donate for it? A second aspect considers where such donations could have the greatest impact. Since the potential benefit per funding unit is usually greatest in developing countries where also need is greatest, most effective altruists donate to development programmes in these countries. Cost-benefit analysis also plays a role in programme selection. Donors decide whether a desired outcome can be achieved more efficiently through, for instance, a health measure or through an education programme. One way to apply EA is to quantify and compare human well-being through a metric called ‘quality-adjusted life years’ (QALY), which seeks to determine, for instance, to what extent a medical treatment increases the life expectancy or the quality of life of a person, compared to a literacy or education programme. It is used to compare the cost-effectiveness of different interventions.³

Origins and proponents of Effective Altruism

The philosophical basis of EA lies in the book *Famine, Affluence, Morality* (1972) by Australian philosopher Peter Singer, a professor of bioethics at Princeton University and representative of classical utilitarianism. Initially regarded as a major thinker behind the international animal rights movement, he argues with regard to humans that all human lives are equal and that people therefore have the same moral obligation to strangers far away as to those close by. Based on this assumption, people should direct their altruistic efforts wherever they will do most good.

William MacAskill, philosophy professor at University of Oxford, is another key figure in the EA community, having published, inter alia, the book *Doing Good Better—Effective Altruism and a Radical Way to Make a Difference* (2015).

These and other advocates regard themselves as members of a new “social movement”,⁴ one that avoids governments and funds programmes or organizations directly. Local groups are active at universities such as Oxford, San Francisco, Melbourne and Berlin. Prominent proponents of the concept can also be found in the Silicon Valley, including the Tesla and SpaceX CEO Elon Musk, Google.org director Jacqueline Fuller, and PayPal and Palantir co-founder Peter Thiel.⁵

The approach is used by a variety of actors, including the so-called Foundation for Effective Altruism, based in Basel, Berlin and San Francisco, which promotes the ideas of Effective Altruism by organizing public events, issuing publications, founding local groups and providing individual advice on how to spend most effectively. Similar activities are undertaken by the Centre for Effective Altruism based in Oxford and Berkeley whose trustees are William MacAskill and Toby Ord.

Giving What We Can, an initiative started by Toby Ord and MacAskill, has built a community of people donating 10 percent of their income to the supposedly most effective organizations.

Another entity using the EA approach is the Open Philanthropy Project, founded in 2017 through a partnership between GiveWell and GoodVentures, a non-profit group set up by Dustin Moskovitz, co-founder of Facebook, and his wife, Cari Tuna. In trying to identify the most effective grant opportunities, the Open Philanthropy Project states: “We believe economic development and technological innovation have greatly increased human well-being. We’re optimistic that this trend will continue, and we hope to play a part in accelerating it.”⁶

There are also consultancy firms or evaluators that advise philanthropists on how to spend their resources most efficiently. One, PHINEO, received, for instance, a 100 million Euro grant from auto maker BMW’s major shareholder Susanne Klatten to support a variety of organizations in Germany. The so-called Skala-Initiative, provides grants only to organizations that have proven to have a social impact.⁷

2 See <https://www.nytimes.com/2015/08/16/upshot/effective-altruism-where-charity-and-rationality-meet.html> and MacAskill (2015).

3 See: <https://concepts.effectivealtruism.org/concepts/measuring-healthy-life-years/>

4 See: <https://ea-foundation.org/effective-altruism/>

5 See: https://www.huffingtonpost.com/melody-y-guan/elon-musk-superintelligen_b_7929960.html?guccounter=1

6 See: <https://www.openphilanthropy.org/about/vision-and-values>

7 See: <https://www.phineo.org/themen/skala-initiative>

One of the most prominent of these consultancy firms is GiveWell, a non-profit group based in San Francisco. Founded in 2007 by Holden Karnofsky and Elie Hassenfeld, two former hedge-fund analysts, the firm recommends a number of recipient organizations that, according to their analysis, use the funds they receive in the most effective way and with a high impact and would thus be trustworthy recipients.

For several years already charities have been rated according to the ratio between their overheads and the amount of funds transferred to beneficiaries. While this way of assessing NGOs has been criticized as not taking sufficiently into account the quality of programme work and being based on a mechanistic accounting, promoters of EA go even one step further towards the economization of development aid. GiveWell, for instance, calculates standardized returns on investment across charities, as measured by criteria such as cost per life saved.

GiveWell evaluates organizations on the basis of four criteria:

- » First, the effectiveness ('Evidence of Effectiveness') of the implemented programme. GiveWell searches for programme activities that have been proven effective through its own assessment as well as by programme beneficiaries and academia.⁸
- » Second, the cost-effectiveness of the implemented programme. Cost-effectiveness according to GiveWell's understanding means "saving or improving lives as much as possible for as little money as possible". According to GiveWell, this is the case for a series of health interventions such as deworming measures and distribution of mosquito nets as well as for cash transfers. GiveWell estimates figures like total "cost per life saved" or "cost per total economic benefit to others".
- » Third, the amount of "room for more funding". GiveWell assesses the added value of additional funds—beyond the resources the organization raises without the support of GiveWell.
- » Finally, the organization's ambition for transparent communication. It must be willing to share comprehensive information about its effectiveness both with GiveWell and with the public.

8 See: <https://blog.givewell.org/2012/08/17/our-principles-for-assessing-evidence/> and <https://blog.givewell.org/2012/08/23/how-we-evaluate-a-study/>

An attractive approach for modern philanthropic foundations

The adoption of EA or elements of it can be observed especially among newly set up philanthropic entities. The new generation of philanthropic initiatives, including that set up by Mark Zuckerberg and his wife Priscilla Chan,⁹ aspires to be more efficient than more traditional foundations. So-called 'venture philanthropies' using business tools such as social impact investment and market-based approaches like public private partnerships to increase the efficiency of charitable interventions.

The Bill & Melinda Gates Foundation is one of the most prominent proponents of Effective Altruism.¹⁰ In their foreword to the 2016 edition of Peter Singer's *Famine, Affluence, and Morality*, Bill and Melinda Gates state that Singer's work is as relevant now as it has ever been. In turn, Singer praises Bill Gates and Warren Buffett for being the "most effective altruists" in history.¹¹

Sue Desmond-Hellmann, CEO of the Bill & Melinda Gates Foundation, reviewed MacAskill's book *Doing good better*, concluding:

"Getting in the habit of giving back is never a bad thing. While good intentions might not be enough in their own right, a world full of people who care—and who are open to doing good better—can make a world of difference."¹² Also, Bill Gates describes MacAskill with the words: "A data nerd after my own heart ..."¹³

For the Gates Foundation, optimizing resources for maximum impact is highly important.¹⁴ In a 2013 article for *Wired* magazine, Bill Gates states:

"I have been sharing my idea of catalytic philanthropy for a while now. It works a lot like the private markets: You invest for big returns. [...] You're working in a global economy worth tens of trillions of dollars, so any philanthropic effort is relatively small. If you want to have a big impact, you

9 See: <https://www.theguardian.com/technology/2015/dec/01/mark-zuckerberg-and-priscilla-chan-announce-new-baby-and-massive-charity-initiative>

10 See: <https://www.gatesfoundation.org/Who-We-Are/General-Information/Foundation-Factsheet>

11 See: <https://www.npr.org/templates/transcript/transcript.php?storyId=529958027&t=1531234328022>

12 See: <https://medium.com/bill-melinda-gates-foundation/the-case-for-putting-your-head-where-your-heart-is-e5523da22f50>

13 See: <https://twitter.com/billgates/status/614046742426845184>

14 See: <https://www.gatesfoundation.org/How-We-Work/General-Information/Evaluation-Policy>

need a leverage point—a way to put in a dollar of funding or an hour of effort and benefit society by a hundred or a thousand times as much. [...] Children were dying of measles for lack of a vaccine that cost less than 25 cents, which meant there was a big opportunity to save a lot of lives relatively cheaply. The same was true of malaria.”¹⁵

The Gates Foundation considers vaccines to be one of the most cost-efficient tools for improving people’s lives. In his annual letter of 2011, Bill Gates talks about the “magic” of vaccines:

“Vaccines have taken us to the threshold of eradicating polio. They are the most effective and cost-effective health tool ever invented. I like to say vaccines are a miracle. Just a few doses of vaccine can protect a child from debilitating and deadly diseases for a lifetime. And most vaccines are extremely inexpensive. For example, the polio vaccine costs 13 cents a dose.”¹⁶

For a long time, the Gates Foundation focused on vertical interventions,¹⁷ mostly narrowly focused technical approaches that promised to be cost-effective and to show quick results. Only after the Ebola outbreak, which exposed the lack of disease preparedness by the international health community, did Bill Gates adjust his opinion, acknowledging the importance of strengthening entire health systems.

In their 2018 annual letter, Bill and Melinda Gates explain, that they focus their grant-making to developing countries according to where the money would have a big impact:

“We don’t compare different people’s suffering. All suffering is a terrible tragedy. We do, however, assess our ability to help prevent different kinds of suffering. When we studied the global health landscape, we realized that our resources could have a disproportionate impact. We knew we could help save literally millions of lives. So that’s what we’ve tried to do.”¹⁸

15 See: <https://www.wired.com/2013/11/bill-gates-wired-essay/>

16 See: <https://www.gatesnotes.com/2017-Annual-Letter>

17 A vertical approach, as in health, is a disease-specific top-down approach with clear measurable objectives in the short or medium term. A horizontal approach is a more comprehensive approach that seeks to treat all the underlying causes for health problems of a population, but is more difficult to measure and shows results only in the longer term.

18 See: <https://www.gatesnotes.com/2018-Annual-Letter>

Risks and side-effects of Effective Altruism

Several concerns go along with the underlying normative assumptions, the methodology and the consequences of the application of EA in practice:

1. Perpetuating the concept of charity

Fundamental to the EA approach is a belief in charity. According to proponents of EA, one should pursue a high-earning career for the purpose of donating a significant portion of one’s income. Following this idea, MacAskill has launched the vocational counseling website “80,000 Hours” that recommends to those starting a career one which could have a high social impact. Central to this is the idea of ‘earning to give’, which could justify even a career in a morally questionable, but high-income industry. The question of ‘where does the money come from?’ is of no concern. Does the wealth earned by business practices that themselves contribute to people’s suffering, as through tolerating exploitative working conditions, or environmentally polluting practices, through investments in weapon companies or food commodity speculation, provide a morally justifiable basis for ‘good’ philanthropy?

The EA approach shares some of the same problems as those of traditional development assistance, including aid dependency and the maintenance of power imbalances, which have been well outlined by Zambian economist Dambisa Moyo among others.¹⁹ In order to end poverty and its related effects globally, neither donations by rich individuals nor development aid by governments will be enough. Needed are fundamental changes in the current socio-economic system, such as access to good health systems and education for all, obliging corporations to respect human rights and environmental standards in their business activities or stopping tax avoidance. Rich and powerful countries that have been benefiting from globalization— at the expense of the less fortunate—have a special responsibility to drive such political and economic changes.

But most important is the need to address existing inequalities. EA does not provide the right concepts to tackle this global challenge as it is grounded in the idea of charity, which in turn requires the

19 See Dambisa Moyo, *Dead Aid: Why aid is not working and how there is another way for Africa* (London: Penguin, 2010).

continuation of wealth accumulation amongst wealthy philanthropists.²⁰ It is unclear whether these so-called ‘high net-worth individuals’ have an intrinsic interest in challenging the existing socio-economic system that allowed them to accumulate their wealth and thereby their ability to give or whether they are merely interested in mitigating the symptoms of inequality.

2. Privileging measurable interventions and technical solutions

Trying to economically maximize the ‘greatest good’²¹ incurs several risks. One is privileging problems that are most easily quantifiable over those that are more difficult to measure. Another is ignoring interventions that can show results only in the longer term. For instance, the value of deworming programmes is measurable both in monetary output and in the value added for the well-being of society and the economy. But how can the value of more complex but nevertheless equally attainable goals like the fulfilment of civil and political rights, of women’s rights, equality, or democracy—all of which contribute to greater health and well-being—ever be calculated in economic terms?

This also raises the question of whether the concept of the ‘greatest good’ is applicable to funding decisions. Is it better in the short term to cure 100 people through a medical intervention or to improve long-term living conditions so that people as a whole less often get sick?

In an interview for the German newspaper *Die Zeit* of July 2016, Melinda Gates explains:

“Bill and I have an economic approach. We ask, where do diseases rage especially? Where do most adults and children die, and what makes their lives especially difficult? There is also a recognized measure for this. With it we find the places in the world where death and disability are particularly present.”²²

20 Among countries, there is a debate on the question of who should be responsible for the provision of public goods and services—the state or wealthy individuals. While some countries have developed strong social welfare states, others rely on private provision of and philanthropic donations for such goods and services. See e.g., Gösta Esping-Andersen, *The three worlds of welfare capitalism* (Cambridge, UK: Polity Press, 1990).

21 A single definition of what EA means by “the good” is not formulated. It mostly refers to helping end people’s suffering, prolong a human’s life or to contribute to changing the world positively (see also Singer 2016, p. 22). An intervention that is considered as most “effective” is then defined as the “greatest good”.

22 See: www.zeit.de/2016/32/melinda-gates-bill-gates-spenden/seite-2

3. Shifting of public resources away from long-term activities

While short-term vertical interventions,²³ such as deworming, and efforts to improve long-term living conditions through profound changes in the underlying social and economic structures both have some merit, the Gates Foundation’s support for vertical interventions has been undermining, directly or indirectly, more systemic approaches. This is particularly evident with regard to health policy, primarily by prompting governments to shift their spending priorities, for instance through so called ‘matching funds’.²⁴ This has become most obvious in the case of the WHO: Due to changing funding patterns of traditional donors, the WHO’s ability to respond adequately to global health emergencies was seriously weakened, as seen in the case of its response to the Ebola outbreak in 2014.²⁵ The influence of big philanthropic foundations—which themselves most often benefit from tax relief—over the use of public resources also undermines democratic control of these resources.

4. Neglecting the recipient voice

One central outcome of the debate on Aid Effectiveness was the recognition by donors that fruitful development cooperation can only be ensured in working together on equal footing with the recipients, be they governments, communities or individuals. In the Paris Declaration on Aid Effectiveness, governments therefore formulated principles of “Ownership”, “Alignment” and “Mutual Accountability” among others:

- » “Ownership: Developing countries set their own strategies for poverty reduction, improve their institutions and tackle corruption.
- » “Alignment: Donor countries align behind these objectives and use local systems.
- » “Harmonisation: Donor countries coordinate, simplify procedures and share information to avoid duplication.

23 Vertical health interventions are disease-specific while horizontal health interventions focus improving general health services, providing prevention and care for prevailing health problems.

24 See Kathrin Hartmann, Interview with medical expert McCoy: The Gates Foundation is a means of exerting power, in *Spiegel Online*, Hamburg, 27 July, 2014; www.spiegel.de/wirtschaft/unternehmen/interview-zur-gates-stiftung-mccoy-beklagt-machtmissbrauch-a-981842.html

25 See also Martens/Seitz (2015).

- » “Results: Developing countries and donors shift focus to development results and results get measured.
- » “Mutual accountability: Donors and partners are accountable for development results.”²⁶

By contrast, the concept of Effective Altruism is, so far, very donor-centrist. Its supporters can be found mostly among people in countries of the global North that have money to make large-scale donations. The measurement of the effectiveness of an intervention and the decision about the choice of which charitable organization should benefit from donations are made by actors from these countries. Unlike the effort by many civil society organizations to apply a participatory and inclusive approach to decision making, EA makes no effort to take into account the recipient voice.

5. The value of a human life as business case

There are several problems related to the methods used by EA consultant firms, such as GiveWell, to calculate the effectiveness of an intervention. EA supporters claim to not only donate more funds than done by other sources of aid, but to use them more effectively by applying business tools. This claim goes along with the belief that business is better placed to solve urgent global social and environmental challenges than are states or conventional philanthropies.²⁷ Donations are treated as investments and should achieve a (measurable) added social value or a ‘social return on investment’. For this purpose, the saving or increased quality of a human life is calculated in economic terms, and the decision about interventions to support made in terms of cost-efficiency. According to many of EA supporters, vaccines are one of the most cost-efficient tools. In his annual letter of 2017, Bill and Melinda Gates explain: “And for every dollar spent on childhood immunizations, you get \$44 in economic benefits.”²⁸

Reading this, the question arises what is the aim of the intervention and what are the means to achieve that end? As the human right to a healthy life does not seem to be sufficient to justify increasing action in this area, economic justifications are needed.

Funding for the provision of medical treatment and ultimately for people’s health is presented as providing good returns to investment, that is, as simply good business.

Although Bill and Melinda Gates deny comparing different people’s suffering,²⁹ this comparison is unavoidable when considering investments in people’s well-being as business case.

By quantifying a human life in monetary terms and defining its economic benefits for society, people become exchangeable and one could argue that the loss of a human’s life at one place could be compensated by saving a different human’s life in another place. A shift in the value canon takes place from the value-in-itself of a human being to the relative cost of ensuring his or her health and well-being.

Conclusion

The underlying question of EA is, without a doubt, a very valid one: How can we make sure that with what people give to charitable organizations, a maximum of good is being achieved, without regard to where recipients of that organization is located, beyond where they can most benefit. In other words, it aims at cleansing charity from selective considerations that should never guide decisions on where to spend scarce resources.

The problems with the concept stem from three fallacies that are embedded in its origins and that will be very difficult to overcome: (a) EA is yet another concept of development emanating from within liberal and elitist circles in the global North, neglecting the recipient voice; (b) it risks diverting attention and—more importantly—funds from urgent issues and other worthwhile endeavors that are by its own definition beyond its scope; and (c) it is to a large degree blinded by short-termism.

The assumptions and implications of EA application contribute towards pushing states and governments further out of their role as the main parties responsible for the provision of public goods and services, while private actors, whether companies or philanthropic donors, are increasingly promoted as more effective in deciding who best to provide such services. This however risks the result that human rights and the institutions and systems

26 See: <https://www.oecd.org/dac/effectiveness/parisdeclarationandaccraagendaforaction.htm>

27 See Matthew Bishop and Michael Green, “Philanthrocapitalism Rising,” in *Society* 52 (2015), pp. 541ff.

28 See: <https://www.gatesnotes.com/2017-Annual-Letter>

29 See: <https://www.gatesnotes.com/2018-Annual-Letter>

which are built on them are increasingly sidelined, undermined or even discarded entirely.

It would be too easy to simply disregard EA as yet another case of something well intentioned, but badly implemented. The questions EA asks can be of great value to anyone trying to determine where or how give. However, nobody should be so foolish to let those decisions be made by algorithms or other quantifiable estimates, however elaborate. At the least, people should be cautious about letting their value systems be marginalized by a practice that appears rational and scientific, but will naturally also be guided by market-based normative assumptions that underlie a veil of numbers and statistics. Apparent contradictions are visible, for example, in the 'earning to give' approach: How can anyone argue that it is ok to choose a high-earning job which risks causing suffering for hundreds of people, just to make enough money to help alleviate some of that suffering?

The key question here, as with many other systems of thought or theory-based models, is scale. What may be the right questions to ask and methods to use for an individual to make giving decisions, can be extremely disruptive when projected onto schemes that by their sheer size become systemic. This is the main message to take away from this analysis: Even when agreeing with the conclusions of EA in terms of cost-effectiveness, they must not become guiding principles for philanthropic donors and policy makers, who should seek ways to address problems over the long term and in a sustainable way. States have to fulfil their responsibility to protect and promote human rights and ensure their realization. Governments have to implement fundamental fiscal and regulatory policies to advance this goal, including regulating business to respect human rights in all their activities and ensure the provision of basic public services and infrastructure, whether in health, education, transport or social protection.

For further reading

Gates, Bill (2013): Here's My Plan to Improve Our World— And How You Can Help. Published online on 11.12.2013 in Wired.

<https://www.wired.com/2013/11/bill-gates-wired-essay/>

Martens, Jens/Seitz, Karolin (2015): Philanthropic Power and Development. Who shapes the agenda?

Aachen/Berlin/Bonn: Brot für die Welt/Global Policy Forum/MISEREOR.

https://www.globalpolicy.org/images/pdfs/GPFEurope/Philanthropic_Power_online.pdf

MacAskill, William (2015): Doing Good Better: How Effective Altruism Can Help You Help Others, Do Work that Matters, and Make Smarter Choices about Giving Back. Avery.

Singer, Peter (2016): The Most Good You Can Do: How Effective Altruism Is Changing Ideas About Living Ethically. Yale University Press.

Singer, Peter (1972): Famine, Affluence, and Morality. In: Philosophy and Public Affairs, Vol. 1, No. 3 (Spring, 1972), pp. 229–243.

http://personal.lse.ac.uk/robert49/teaching/mm/articles/Singer_1972Famine.pdf

Iason, Gabriel (2016): Effective Altruism and its Critics. In: Journal of Applied Philosophy, Vol.34, No.4, 457–473.

<http://onlinelibrary.wiley.com/doi/10.1111/japp.12176/pdf>

Imprint

What you get for your dollar

“Effective Altruism” — What it is, how philanthropic foundations use it and what are its risks and side-effects

Published by

Global Policy Forum Europe e.V.
Königstr. 37 a
53115 Bonn
Germany

Tel. +49-228-96 50 510
europe@globalpolicy.org
Contact: Karolin Seitz

Evangelisches Werk für Diakonie und Entwicklung
Brot für die Welt
Caroline-Michaelis-Str. 1
10115 Berlin
Germany

info@brot-fuer-die-welt.de
www.brot-fuer-die-welt.de
Contact: Mareike Haase

Bischöfliches Hilfswerk MISEREOR
Mozartstr. 9
52064 Aachen
Germany

info@misereor.de
www.misereor.de
Contact: Klaus Schilder

Author: Karolin Seitz

Editorial Assistance: Laura Angresius, Mareike Haase, Carolin Hornbach, Karen Judd, Wolfgang Obenland, Klaus Schilder, Stig Tanzmann

Design/Production: www.kalinski.media

Aachen/Berlin/Bonn, January 2019