

# Effective Altruism

---

## Introduction

The world we live in today is home to many serious problems. You might have heard that over 750 million people live on less than US \$2 per day<sup>1</sup>, or that millions of children die each year of easily preventable causes such as malaria, diarrhea, or pneumonia.<sup>2</sup> You know that climate change threatens environmental havoc and is set to cost the economy trillions of dollars<sup>3</sup>. Perhaps you've read that, despite social progress, a third of women worldwide have suffered from sexual or other physical violence in their lives<sup>4</sup>, and that the number of people living with depression has increased substantially over the last decade<sup>5</sup>. On top of all this, existential risks such as a global pandemic or nuclear war could threaten the fate of human civilisation if we fail to do something about them now.

Given the severity of these problems, surely we have a responsibility to do something about them. Many of us, when faced with suffering, injustice and death, are moved to want to make a difference. But what can we do? There are countless problems and many different possible ways of addressing them. In addition, as individuals and globally, we don't have the luxury of unlimited resources which would allow us to solve all these problems at once. We have to make decisions about how to allocate the resources we have, in order to have the greatest impact. Just as a nurse must triage patients in order to ensure that the emergency doctor can help the maximum number of people possible, we too should base our decisions about helping the world on evidence that leads us to having the greatest impact possible.

A growing social movement has pioneered one approach to doing this, and is rapidly gathering steam as a community consisting of thousands of people around the world who have decided to use their time on earth to improve the world as much as they can. This is the *effective altruism* movement, and it promotes a moral framework and methodological approach to prioritizing causes, charities and even careers in order to work out how we can have the greatest impact possible, and increase the chances of creating a flourishing world for future generations. Effective altruism is about using evidence and reason to figure out how to

---

<sup>1</sup> (at purchasing power parity) World Bank Group, *Poverty and Shared Prosperity: Taking on Inequality 2016* (Herndon: World Bank Publications, 2016), ch. 2

<sup>2</sup> UNICEF, "Levels & Trends in Child Mortality", [https://www.unicef.org/publications/files/Child\\_Mortality\\_Report\\_2017.pdf](https://www.unicef.org/publications/files/Child_Mortality_Report_2017.pdf).

<sup>3</sup> John Broome, *Climate Matters: Ethics in a Warming World* (New York: W.W. Norton, 2012); William Nordhaus, *The Climate Casino: Risk, Uncertainty, and Economics for a Warming World* (New Haven: Yale University Press, 2015).

<sup>4</sup> United Nations, *The World's Women 2015: Trends and Statistics* (New York: United Nations, 2015).

<sup>5</sup> World Health Organization. "Depression and other common mental disorders: global health estimates." (2017).

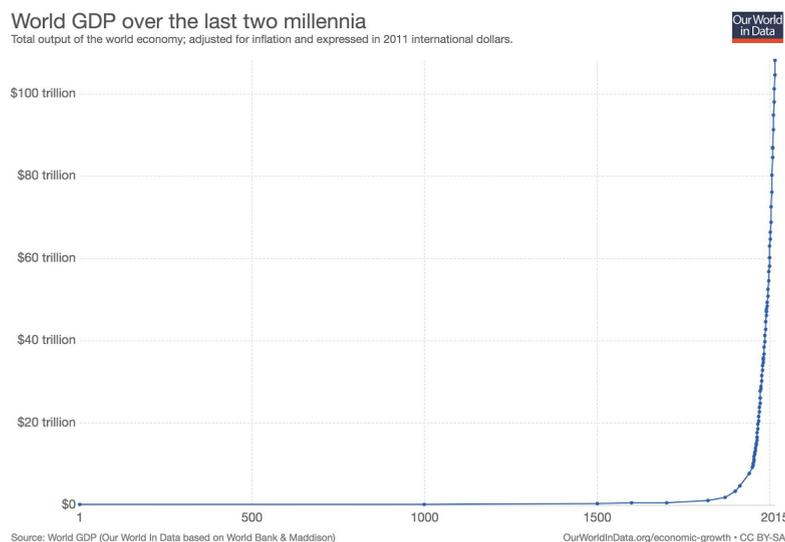
help others as much as possible, and taking action on that basis. In this chapter, I'll explain why effective altruism is so important in the world today, introduce the framework that effective altruism uses to respond to the challenge of knowing which problems to work on, and illustrate how it can help us to create a flourishing world.

## We're at an unique point in history

Reading the list of problems with which I opened this chapter, you'd be forgiven for feeling a sense of hopelessness—perhaps there's no way for any one of us to stop the decline of world progress. That'd be a huge shame, because in that sentiment, we'd be completely wrong. We're at a particularly unique point in human history where we have access to more resources than we've ever had before, and therefore more opportunity to have a huge positive impact on some of the world's most urgent and important problems.

If you were to look at a graph that represents the entire economic history of human civilisation, you'd be looking mostly at a fairly flat line, with not much going on. That's because, for the vast majority of human history, just about everybody lived on the equivalent of \$1 per day, and not much changed for thousands of years.

But then about 250 years ago, something extraordinary happened: the scientific and industrial revolutions. Suddenly, economic development skyrocketed, resulting in more growth over the next couple of hundred years than human civilization had seen in the thousand years before. Human society experienced an unprecedented increase in available resources which led to tremendous economic and technological progress.



What's more, this unusual period of accelerated progress has left a great number of us extraordinarily privileged, able to enjoy the kind of comfortable existence facilitated by

advancements in technology, medicine and entertainment that our great-grandparents couldn't have dreamed of.

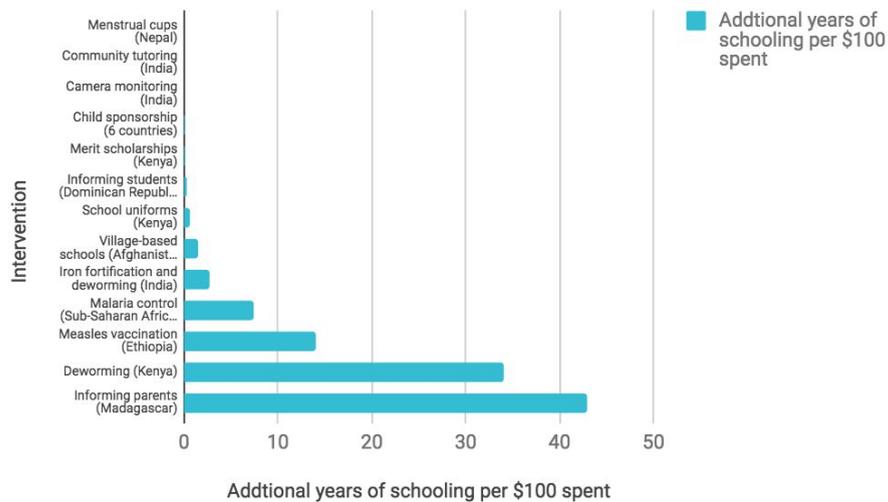
So how should we use our tremendous economic and technological power to improve the world? There are a number of things we need to address in order to answer this question, such as whether to do good through your career, donating to charity, or political engagement; what causes to focus on and who you should choose to work with. But the most fundamental question is, what problems should we be trying to solve if we want to have the greatest positive impact possible? We must make decisions about how to allocate the resources we have. These decisions are incredibly important, because though our available resources are greater than ever, they're still finite. We don't have all the time in the world to dedicate to solving every single problem. The choices we make about what to focus on matter, as we can do far more good by focusing on some problems than on others.

### **Some ways of doing good are vastly better than others**

With many choices in life, most of us understand that it's important to base our decisions on evidence and reason rather than guesswork or gut instinct. When you make a decision about spending a lot of money, such as buying a car or booking holiday accommodation, you usually do a bit of research in order to find a good deal; for example reading customer reviews or asking friends for recommendations. Certainly, you'd feel cheated if you spent 100 times more on a car, phone, or hotel room than you needed to. Yet when it comes to focusing on global problems, many of us are far less discerning. This is a problem, as the best ways of helping others are often vastly better than others.

Below is a chart comparing the effectiveness of different interventions on school attendance rate (how many additional years of schooling did each intervention provide per \$100 spent on that intervention). The chart shows several different strategies including providing menstrual cups to girls in Nepal, merit scholarships and school uniforms in Kenya, iron fortification in India, and informing parents of the benefits of education in Madagascar.

## Additional years of schooling per \$100 spent vs Intervention



As the above graph illustrates, the best interventions (deworming programs in Kenya and providing parents of school children with information about the benefits of education in Madagascar) have a far larger effect on getting children to stay in school than programs which provide things like menstrual cups, school uniforms or even child sponsorship programs. It might seem cold and unfeeling to make trade-offs in this way, particularly when we're discussing the lives and well-being of real human beings. But that is precisely why we have a moral obligation to make these trade-offs. Whether you like it or not, every time you choose to donate to a particular charity or work on a certain problem, you're making an implicit decision *not* to donate somewhere else, or a decision *not* to focus your attention on other problems.

The results of research into common charitable interventions are often counterintuitive and surprising. For example, charities that distribute school supplies to schools in poor countries actually have very little impact on school retention rate (getting kids come to and stay at school), or their test scores. In other words, a \$100 donation to these types of charities would probably do very little to improve education levels amongst children in the developing world. However, for the same amount of money, by educating mothers about the educational benefits of schooling, you could help children gain over 40 years of additional education.

These large differences between intervention effectiveness are likely to be just as large in other cause areas, but these are often very difficult to tell apart without having high-quality research and evidence or knowing which experts to trust. Comparing the different ways of doing good is difficult, both emotionally and practically, but these comparisons are vital to ensure we do as much good as possible. Which interventions have the highest impact is an important open question, one the effective altruism movement is trying to answer.

## What is Effective Altruism?

Effective Altruism is a philosophy and research project that attempts to answer the question: *how can we do the most good?* It is practical as well as philosophical, and there a community of people taking these questions seriously and focusing their efforts on finding the most promising solutions to the world's biggest problems.

There are some defining characteristics of effective altruism. These are:

- *Maximising.* Effective Altruism aims to do *as much good* as possible, not just some good, or enough good.
- *Science-aligned.* The scientific method is the best means we have to figuring out how to do the most good. Broadly construed, this means relying on empirical observation, careful rigorous argument and theoretical models.
- *Tentatively welfarist.* Effective Altruism holds as a tentative hypothesis that goodness means improving the welfare of individuals.
- *Impartial.* Everyone's welfare counts equally, regardless of nationality or location.

### **Note: Effective Altruism is not Utilitarianism**

It is important to distinguish effective altruism from utilitarianism. Effective altruism shares similarities with utilitarianism: it's about maximising good and improving well-being, and reducing suffering as much as possible. It's also true that members of the community choose to make significant personal sacrifices (such as donating a large percentage of their income) in order to do more good. However, unlike utilitarianism, effective altruism does not claim that we should sacrifice our own interests if we can benefit others to a greater extent. Nor does effective altruism claim that all ways of helping others are morally permissible as long as they help others the most; in fact the effective altruism community has strong norms against engaging in any activities which may cause harm. Finally, effective altruism (and all plausible moral views), recognize that well-being is at least part of the good that we should be looking to maximise.

### **A framework for solving the world's biggest problems**

Thinking about which problems to solve, or *cause prioritization*, is incredibly important, because the impact of focusing on one cause over another could differ by orders of magnitude.

One might be inclined to focus just on whatever problems are *biggest*. But that's not the only factor to consider. A problem may be very big, but also incredibly difficult to solve: for example, it would be great if we could invent a perpetual motion machine to supply us with free energy forever, but we're fairly confident that this is impossible. So it's not worth pouring endless amounts of time and energy into this project. What's more, some causes aren't worth focusing on because others have already taken the best opportunities to make progress, and adding additional resources won't make much of a difference. For example, cancer research is important as cancer affects huge numbers of people and causes a lot of suffering; but governments and corporations already spend millions funding cancer research, and it's unlikely that you'd be able to have much of an impact on this problem by providing additional funding. The question is how we weigh up different causes in order to find those for which we'll have as much impact as possible by working on them.

Effective Altruism gives us a framework for answering that question, and this framework is actually very simple: it says that a problem is higher priority the bigger, the more easily solvable, and the more neglected it is.

### **The cause-prioritization framework**

The most commonly used heuristic in the effective altruism community to figure out which causes to focus on is a three-factor cause-prioritization framework. Using this framework, we can work out how important a cause is by weighing up the following factors:

#### ***Scale***

The number of people affected and the degree to which they are affected.

#### ***Solvability***

The fraction of the problem solved by increasing the resources by a given amount. (This factor is also referred to as *tractability*.)

#### ***Neglectedness***

The amount of resources already going toward solving the problem

Using the framework above, effective altruists have come up with a list of causes that constitute our best guesses of what are the most effective causes to work on. These are currently some of the most promising cause areas, but could change significantly over time as

we gather more evidence through research.

## **Farm Animal Welfare**

60 billion animals are killed for food each year, and the vast majority of these spend their lives in factory farms in horrendous conditions; crammed together with little space, natural light or stimuli, and at constant risk of developing problems such as weakened or broken bones, infections and organ failure. Most have their lives ended prematurely when they are slaughtered for food.

Farm animal welfare is highly *neglected*. Despite the size of the problem, in the US only a few tens of millions of philanthropic dollars are donated every year to organizations that focus on improving the lives of farmed animals—0.01% of total US philanthropy. The amount is tiny even compared to other animal causes: there are 3000 times more animals in factory farms than there are stray pets, but factory farming gets one-fiftieth of the funding. For this reason, additional funding can have a huge impact in this area, for example by enabling activists to campaign for large retailers and fast food chains to cut caged eggs out of their supply chains. Because of the sheer numbers of sentient beings involved, making progress on improving farm animal welfare could avert a huge amount of suffering.

## **Global health and development**

Areas within global health and development such as alleviating extreme poverty and fighting infectious diseases are considered a high priority within the effective altruism community, as they're particularly *solvable*. That is, efforts in global health have a great track record of making an impact. For example, the eradication of smallpox in 1973 saved over 60 million lives, and there is a large body of high-quality scientific evidence, including randomized controlled trials and meta-analyses, assessing the effectiveness of different programs.

Many interventions in these areas are incredibly cost-effective, too. Antimalarial bednets cost around \$5.00 each. GiveWell, an independent charity evaluator, estimates that they can significantly reduce malaria rates.<sup>6</sup> Even simply transferring money to people who are very poor is a relatively cost-effective way of helping people.<sup>7</sup> Other top-recommended interventions include deworming schoolchildren, seasonal malaria chemoprevention, and vitamin A fortification.

---

<sup>6</sup> "Mass Distribution of Long-Lasting Insecticide-Treated Nets (LLINs)" *GiveWell*, November 2013, Last updated March 2018.

<https://www.givewell.org/international/technical/programs/insecticide-treated-nets>

<sup>7</sup> "Economic Empowerment" *Giving What We Can*, April 2016.

<https://www.givingwhatwecan.org/report/cash-transfers/>

## **Existential Risks**

An *existential risk* is a risk of an event that would either annihilate intelligent life on Earth or permanently and drastically curtail its potential. These, though low in probability, are extremely great in *scale*.

To see this, consider how bad an existential catastrophe would be. It would involve the deaths of all 7.5 billion people on this planet, including you and everyone you've ever known or loved. But it would also mean the end of humanity's story: the loss of all of our future potential. That's a tragedy on an unimaginable scale. There's no reason why, if we're careful, humanity couldn't survive for many millions of years. The extinction of the human race would therefore mean the preclusion of *trillions* of lives to come.

What's more, the future could be very good indeed. Despite there being a lot of very big problems in the world, there's plenty of reason to believe that the world is getting better over time. Over the last 200 years we've seen incredible progress in the increase of people's life expectancy, the reduction of global poverty and progress made in social issues such as women's liberation, the decriminalisation of homosexuality and the spread of democracy. If we have continued progress, the heights of human achievement we might reach could be unimaginably great.

Finally, there are concrete things we can do about existential risk. Rapid technological progress has brought with it tremendous prosperity and the opportunity to do, create and discover great things, but that same technological progress has also brought greater risks, such as the risk of nuclear war and the possibility of extreme climate change. Other existential risks include bioterrorism, via the creation of a highly contagious and lethal virus, or the development of an artificial general intelligence with misaligned goals that could accidentally destroy civilisation. Preventing these types of risks may be one of the most important moral issues we face.

### **What does this mean for you?**

Effective altruism is an ongoing research project, and we still don't know how to do the *most* good—there is still much that remains uncertain and unexplored. Ultimately, we need some an ethical revolution, to shift the way people think about doing good, and cause people to feel inspired to accept the responsibility to use their time, resources and careers to help the more in the most effective way possible.

You exist at a particularly unusual point in human history, and you have a tremendous opportunity to make a positive difference to the world. By choosing wisely, over your lifetime you could donate enough money to effective charities to save dozens of human lives,

or spare hundreds of thousands of chickens from immense suffering in confined cages. By dedicating your career to working on the most pressing problems or using your skills to help highly effective organisations, you could increase the chance of preserving the future of civilization for millions of years to come. We have the resources and capacity to help others enormously at very little cost to ourselves. So we should.

## Education graph footnotes:

| Intervention                             | Source  |
|--|---|
| Menstrual cups (Nepal)                   | Oster, Emily, and Rebecca Thornton. "Menstruation, sanitary products, and school attendance: Evidence from a randomized evaluation." <i>American Economic Journal: Applied Economics</i> 3.1 (2011): 91-100.  |
| Community tutoring (India)               | Banerjee, Abhijit V., et al. "Pitfalls of participatory programs: Evidence from a randomized evaluation in education in India." <i>American Economic Journal: Economic Policy</i> 2.1 (2010): 1-30.   |
| Camera monitoring (India)                | Duflo, Esther, Rema Hanna, and Stephen P. Ryan. "Incentives work: Getting teachers to come to school." <i>American Economic Review</i> 102.4 (2012): 1241-78.   |
| Child sponsorship (6 countries)          | Wydick, Bruce, Paul Glewwe, and Laine Rutledge. "Does child sponsorship pay off in adulthood? an international study of impacts on income and wealth." <i>The World Bank Economic Review</i> 31.2 (2016): 434-458.  |
| Merit scholarships (Kenya)               | Kremer, Michael, Edward Miguel, and Rebecca Thornton. "Incentives to learn." <i>The Review of Economics and Statistics</i> 91.3 (2009): 437-456.  |
| Informing students (Dominican Republic)  | Jensen, Robert. "The (perceived) returns to education and the demand for schooling." <i>The Quarterly Journal of Economics</i> 125.2 (2010): 515-548.   |
| School uniforms (Kenya)                  | Evans, David, Michael Kremer, and Muthoni Ngatia. "The impact of distributing school uniforms on children's education in Kenya." (2008).  |
| Village-based schools (Afghanistan)      | Burde, Dana, and Leigh L. Linden. "Bringing education to Afghan girls: A randomized controlled trial of village-based schools." <i>American Economic Journal: Applied Economics</i> 5.3 (2013): 27-40.  |
| Iron-fortification and deworming (India) | Bobonis, Gustavo J., Edward Miguel, and Charu Puri-Sharma. "Anemia and school participation." <i>Journal of Human resources</i> 41.4 (2006): 692-721.   |
| Malaria control (Sub-Saharan Africa)     | Kuecken, Maria, Josselin Thuilliez, and Marie-Anne Valfort. "Does malaria control impact education? A study of the Global Fund in Africa." (2014).  |
| Measles vaccination (Ethiopia)           | Anekwe, Tobenna D., et al. "The causal effect of childhood measles vaccination on educational attainment: A mother fixed-effects study in rural South Africa." <i>Vaccine</i> 33.38 (2015): 5020-5026.  |
| Deworming (Kenya)                        | Miguel, Edward, and Michael Kremer. "Worms: identifying impacts on education and health in the presence of treatment externalities." <i>Econometrica</i> 72.1 (2004): 159-217.<br><br>Baird, Sarah, et al. "Worms at work: Long-run impacts of a child health investment." <i>The quarterly journal of economics</i> 131.4 (2016): 1637-1680. |
| Informing parents (Madagascar)           | Nguyen, Trang. "Information, role models and perceived returns to education: Experimental evidence from Madagascar." <i>Unpublished manuscript</i> 6 (2008).  |