

MASTERARBEIT / MASTER'S THESIS

Titel der Masterarbeit / Title of the Master's Thesis

„Making the world a better place with science and donations?
A critical review of Effective Altruism and its introduction
into ethics class“

verfasst von / submitted by

Paul Jakob Jurtschitsch, Bakk. BEd

angestrebter akademischer Grad / in partial fulfilment of the requirements for the degree of

Master of Education (MEd)

Wien, 2021 / Vienna 2021

Studienkennzahl lt. Studienblatt / degree
programme code as it appears on
the student record sheet:

UA 199 500 525

Studienrichtung lt. Studienblatt /
degree programme as it appears on
the student record sheet:

Masterstudium Lehramt Sek (AB) UF
Bewegung und Sport UF Psychologie
und Philosophie

Betreut von / Supervisor:

Mag. Mag. Mag. Dr. Dr. Paul Tarmann

Acknowledgement

I would like to thank Peter Singer and William MacAskill for trying really hard to make the world a better place and Paul Tarmann for giving me the opportunity to write about this topic.

Table of Contents

I.	<i>List of Figures</i>	7
II.	<i>List of Abbreviations</i>	8
1	<i>Introduction</i>	9
1.1	The structure of this thesis	10
2	<i>The raison d'être of this thesis</i>	11
2.1	Religion in Austria	11
2.2	Religion in Austrian schools	11
2.3	Children with no religious affiliation were long neglected	12
2.4	Mandatory ethics classes at school as of 2021/2022	12
2.5	The contribution of this thesis to the Austrian ethics classes	13
3	<i>Defining Effective Altruism</i>	14
4	<i>The ethical argument for Effective Altruism</i>	17
4.1	In the classroom	19
4.2	The logical argument	20
4.3	In the classroom	21
4.4	Among the 10 percent	22
4.5	In the classroom	24
5	<i>Five key aspects of Effective Altruism</i>	25
5.1	Scale	25
5.1.1	A constant situation of crisis	26
5.1.2	Choosing which charity to support	26
5.1.3	A situation of crisis requires us to set priorities	27
5.1.4	Assessing the impact of health interventions: the QALY	27
5.1.5	Limitations of the QALY	29
5.1.6	The impact gap between health interventions	29
5.1.7	Making the right questions	30
5.2	Effectivity	31
5.2.1	PlayPump: an aid program gone wrong	31
5.2.2	Aid programs changing the world for the better	33
5.2.3	Finding the needle in the haystack with the help of science	33
5.2.4	Improving school attendance in Kenya	34
5.2.5	The best way to help might not be the most intuitive idea	34
5.2.6	Not only effective but extremely effective	35

5.2.7	Treatment with long-lasting positive effects.....	35
5.2.8	A fat-tailed distribution in effectiveness.....	37
5.3	Neglectedness	38
5.3.1	Diminishing marginal utility: the more you have it, the less it is worth	38
5.3.2	Increasing the value of your donation by supporting the neglected causes.....	39
5.3.3	Emotions are not always the best guide	39
5.3.4	The effect of an individual's story	40
5.3.5	Out of sight, out of mind	40
5.3.6	We are already living in a constant situation of crisis	41
5.3.7	Sudden tragic events harness our emotions	41
5.3.8	Your donation is many times more worth overseas	42
5.3.9	Malaria: once a global problem, today the poor's burden	42
5.3.10	Recognizing those low-hanging fruits.....	43
5.4	Alternative.....	44
5.4.1	Someone else might do a better job	44
5.4.2	PlayPump: causing more harm than good	44
5.4.3	What are the alternatives?.....	45
5.4.4	Replaceability	45
5.4.5	Professional philanthropy: an alternative for everyone	46
5.4.6	Among the top 10 percent	47
5.5	Expected value	48
5.5.1	Micromort: assessing risk.....	48
5.5.2	The Fukushima Daiichi nuclear accident: a black swan	49
5.5.3	Expected value: assessing possibilities.....	50
5.5.4	Expected value: comparing potential careers	50
5.5.5	The expected value of political change.....	52
5.5.6	Why should the status quo be the end of political change?	52
5.5.7	Low-probability high-impact risks	53
5.6	In the classroom.....	54
6	<i>Effective Altruism's biggest cause areas</i>	55
6.1	Shifting focus in the Effective Altruism community.....	56
6.2	Global poverty	58
6.2.1	Extreme poverty is the breeding ground for extreme suffering	59
6.2.2	Covid19: the first progress setback in 20 years	60
6.2.3	Poverty hits the most vulnerable segments of the population harder.....	61
6.2.4	Organizations working on this cause area	63
6.3	Reduction of catastrophic and existential risks.....	66
6.3.1	Risks with varying probability	66
6.3.2	The risk of a pandemic	68
6.3.3	Artificial Intelligence.....	71

6.4	Cause prioritization.....	74
6.5	Non-human animals	75
6.5.1	Do non-human animals count?	76
6.5.2	Especially neglected cause	78
6.5.3	Wild-animal suffering	79
6.5.4	Organizations working on their behalf	80
6.6	Other cause areas	81
6.7	In the classroom.....	82
7	<i>Criticism of Effective Altruism</i>	83
7.1	Implicit assumptions within Effective Altruism.....	83
7.1.1	The dominance of consequentialism	84
7.1.2	Welfarism.....	87
7.1.3	The Scientific Approach.....	88
7.2	The Institutional Critique of Effective Altruism.....	89
7.2.1	The problem with earning to give.....	91
7.2.2	Counterargument to the institutional critique	92
7.3	Elitist Philanthropy	93
7.4	In the classroom.....	94
8	<i>Effective Altruism in ethics class.....</i>	95
8.1	The Austrian ethics class curriculum.....	95
8.2	Lesson plans	96
8.2.1	Lesson plan 1: The ethical argument for EA	97
8.2.2	Lesson plan 2: The logical argument for EA	98
8.2.3	Lesson plan 3: Among the 10% & Five key aspects of Effective Altruism.....	99
8.2.4	Lesson plan 4: Effective Altruism's biggest cause areas	100
9	<i>Conclusion.....</i>	101
10	<i>Summary.....</i>	103
11	<i>Abstract.....</i>	104
12	<i>Bibliography</i>	105

I. List of Figures

Figure 1 Global Income Distribution in 2011. Graph adopted from Wiblin (2017).....	23
Figure 2 Health profiles both with intervention i (solid line) and without intervention (broken line). Figure adopted from Sassi (2006, p. 404).	28
Figure 3. A comparison of the size of the effects in additional school years per US\$ 100 invested in different programs. The results are based on the evaluation of different treatments by different studies. Figure adopted from Figure 6 from The World Bank (2012, p. 18).	36
Figure 4 Global income distribution in 2003 and 2013. Graph adopted from Roser (2013)...	47
Figure 5 Comparison between 2018 and 2019 responses in the Effective Altruism community survey when participants were asked to choose one cause to be the top priority in Effective Altruism. Graph adopted from Dullaghan (2020).....	56
Figure 6 Likert graph depicting cause areas listed in descending order based on the average of their full scale ratings from the Effective Altruism survey 2019. Graph adopted from Dullaghan (2020).....	58
Figure 7 World map of Disability-Adjusted Life Years (DALYs) per 100,000 individuals from all causes. DALYs measure the total burden of disease - both from years of life lost due to premature death and years lived with disability. One DALY equals one lost year of healthy live. Figure adopted from Roser & Ritchie (2016)	60
Figure 8 Deaths globally by age in 2017. Graph adopted from Roser et al. (2019).....	61
Figure 9 Chart of the causes of death for childre under 5 years old. The area of each rectangle corresponds to the annual number of deaths of children under 5 years old. Graph adopted from Dadonaite (2019)	62
Figure 10 Number of animals used and killed by humans in contrast to the money donated to animal charities in the U.S.. Figure adopted from Animal Charity Evaluators (2016)	78

II. List of Abbreviations

ACE	Animal Charity Evaluators
AGI	Artificial General Intelligence
AI	Artificial Intelligence
CEA	Centre for Effective Altruism
CSER	Centre for the Study of Existential Risk
DALY	Disease-adjusted Life Years
EA	Effective Altruism
ERS	Existential Risk Studies
FHI	Future of Humanity Institute
FLI	Future of Life Institute
GCRI	Global Catastrophic Risk Institute
GPP	Global Priorities Project
GWWC	Giving What We Can
ICS	Investing in Children and Their Societies
MIRI	Machine Intelligence Research Institute
MERS	Middle East Respiratory Syndrome
NGO	Non-governmental organization
NPO	Non-profit organization
RCTs	Randomized Control Trials
SARS	Severe Acute Respiratory Syndrome
SARS-CoV-1	SARS coronavirus species 1
SKAT	Swiss Resource Centre and Consultancies for Development
TLYCS	The Life You Can Save
QALY	Quality-adjusted Life Years
WHO	World Health Organization

1 Introduction

Looking at our world today, it is easy to get overwhelmed by the vast number of problems we face. From the current pandemic razing with the livelihoods of millions of people, especially the poor in developing countries to the fires raging around the world due to rising global temperatures. Democratic backsliding (Hyde, 2020) and increasing international tensions seem to be on the rise worldwide (National Intelligence Council, 2020). All this, while millions of people suffer from devastating diseases, such as cancer, diabetes, malaria and HIV/AIDS.

With so many pressing problems around the world, how should we decide which problems to tackle first and how should we address them? We are certainly unable to fix all problems at once, given that we only have a limited number of resources. This means we need to make decisions about where and how we want to allocate these resources. The question is on what basis we should make such choices?

This is exactly the question the proponents of Effective Altruism (EA) attempt to address. The aim of EA is, as its name indicates, to incorporate the idea of effectiveness into *altruism*, understood as the “behavior intended to benefit another, even when this action risks possible sacrifice to the welfare of the actor” (Monroe, 2001, p. 415). William MacAskill, one of the founders of the movement proposed an alternative understanding of the concept of altruism excluding the requisite of a sacrificial connotation: “[...] [I]f you can do good while maintaining a comfortable life for yourself, that’s a bonus, and I’m very happy to call that altruism” (MacAskill, 2016, p. 12). As for *effectiveness*, William MacAskill describes it as “[...] doing the most good with whatever resources you have.” (MacAskill, 2016, p. 12). So, “[...] effective altruism is not just about making *a* difference or doing *some* amount of good. It’s about trying to make the *most* difference you can” with whatever resources you have (MacAskill, 2016, p. 12).

Today, the Effective Altruism community entails thousands of people around the world, many of whom have committed to donate a share of their income to charities they deem effective and or have chosen their careers, at least in part, on the basis of EA. Effective altruists – members of the Effective Altruism community – have started non-profit organizations (NPOs), start-ups, and think tanks, they have gone into research, politics and even finance (to donate a big chunk of their earnings) since learning about EA. Each year, thousands of people attend Effective Altruism Global conferences around the world. Over \$1.5 billion of lifetime donations have been pledged by thousands of effective altruists for organizations they consider to be the most

cost-effective. (MacAskill, 2019b, p. 10f). In 2019 alone, GiveWell was able to steer more than \$152 million to its top recommended charities (GiveWell, 2020e). Good Ventures distributes over \$200 million each year in grants to (new) initiatives in line with Effective Altruism values (Good Ventures, n.d.; MacAskill, 2019b, p. 11).

Seeing the impact Effective Altruism is having in the world and the overarching question effective altruists occupy themselves with, I consider the subject worthy of discussion as an interesting topic for ethics classes in high schools. As a future ethics class teacher in the Austrian high school system, I elaborate in this thesis one approach to introduce the subject of Effective Altruism into ethics class in the form of a lesson series. For this, I first present the reader with the main aspects of Effective Altruism with a short discussion in each chapter on how I plan to discuss the subtheme with the students. I also present four possible lesson plans in chapter 8.

1.1 The structure of this thesis

To give the reader a good overview of Effective Altruism, this thesis will be segmented into various parts. The next chapter will discuss the *raison d'être* of this thesis, followed by the definition(s) of Effective Altruism in chapter 3. The fourth chapter will cover the ethical (and logical) argument for Effective Altruism, elaborated by Peter Singer (2009). Chapter 5 elaborates on five key aspects used by effective altruists to determine how to do the 'most good'. The chapter is primarily guided by William MacAskill's *Doing Good Better* (2016) with additional information the author of this thesis found adequate to add.

This overview of the concepts of Effective Altruism is then followed by a short overview of its current main areas of focus, so-called *cause areas*. This gives the reader an overview of the fields of activity of the *Effective Altruism community* – people whose actions have been shaped by the principles of Effective Altruism to a greater or lesser extent. After having presented the concept of Effective Altruism and its main cause areas, chapter 7 covers some of the main criticism of the movement.

As an aspiring ethics teacher, I present in the last part of this thesis four lesson plans I drafted, illustrating how I am planning to introduce the subject into ethics classes. My goal is to engage high school students in ethics class into a discussion about Effective Altruism. This is then followed by the conclusion and the summary of the thesis, as well as the abstract, both in English and German.

2 The raison d'être of this thesis

In Austria, religion classes are mandatory for children with legally recognized religious confessions in primary and secondary school. The big majority of children deregistered from religious class have had no other option than to spend the hour outside the classroom while their fellow students remained in class; merely a few have had the option to attend ethics classes as an alternative. Only until recently did the Austrian government finally decide to set ethics classes as mandatory for those not attending religious classes.

2.1 Religion in Austria

The majority of Austrian residents (64 percent) describe themselves as roman-Catholics. The second biggest religious group are Muslim, with eight percent of the population. Protestants and orthodox make up five percent of the population each. And 17 percent, that is almost as many as Muslims, Protestants and Orthodox combined, are non-denominational, ergo, without a religious affiliation (OIF, 2019, p. 2).

A change of religious confession for the child is only possible with the consent by both parents (or legal guardians). No single parent is allowed to deregister his or her child from religious classes. At the age of twelve, children are able to hinder their parents in changing their religious confession. Only at the age of 14 do children get their own rights to decide which religious confession he or she wants to adopt or whether he or she wants to turn down any religious affiliation. (oesterreich.gv.at, 2020).

2.2 Religion in Austrian schools

According to Austrian law, all children belonging to any recognized religious group are to receive fitting religion classes at school. Attendance to religion classes is mandatory at public schools, which make up the majority of schools in Austria. The classes are monitored and planned by their respective religious associations, who train and choose the religion teachers. This is part of an integration effort reinforced by the Austrian Integration Fund (oesterreich.gv.at, 2020).

Catholic religious classes are available in all schools in Austria. Islam religious classes, covering the second biggest religion group in Austria, were first established in the school year of 1982/83. Today 2,272 schools in Austria offer Islam religion classes (OIF, 2019, p. 4). Religion classes are also available at some schools for Orthodox as well as Protestants.

2.3 Children with no religious affiliation were long neglected

As discussed above, 17 percent of Austrians have no religious affiliation and thus deregistered their children from religion classes. However, until now there has been no replacement class at all and children have had to leave their classrooms while their classmates stay inside. Only a few of them have had the option to attend ethics classes instead, a rare offer in some Austrian high schools. In fact, in 2020 there were only 233 schools of the secondary education level (de. Sekundarstufe II), a tiny fraction, offering ethics classes as a replacement to religion classes in Austria (BMBWF, 2020).

This left several students who are either atheist, agnostic or have a religious faith for which no religion classes are offered in Austria schools, strayed. And this despite there being public demand for the introduction of mandatory ethics classes as an alternative to religion classes to the Austrian National Assembly since 2011 (BMBWF, 2020).

2.4 Mandatory ethics classes at school as of 2021/2022

Only until last year, on 20 November 2020, did the Austrian National Assembly made the decision to adopt the resolution to set ethics classes as mandatory for students not attending religion classes. The plan is to implement this in all schools at the secondary education level from the next school year (2021/22) on. These include both general and academic schools (ge. *Allgemein höhere Schulen*), and higher-level technical and vocational colleges (ge. *berufsbildenden mittleren und höheren Schulen*) (BMBWF, 2020).

The curriculum of the ethics classes in Austria is, according to the Federal Ministry of Education, Science and Research, currently still in development. Its structure shall be based in part in the experimental ethics classes curricula they implemented in the last years with philosophy as its main core. It will also entail various cross-cutting subjects, such as law, history and economics, biology, and political science, among others (BMBWF, 2020).

As for the ethics teachers, higher education institutions have been offering a coordinated curriculum for ethics of 60 ECTS for in-service training and professional development. The University of Graz as well as the University of Vienna offer an ethics-specific master program of 120 ECTS since 2016. A teacher's training in the form of a bachelor degree (160 ECTS) is forecasted to be available at earliest in 2022 (BMBWF, 2020).

2.5 The contribution of this thesis to the Austrian ethics classes

As an aspiring ethics schoolteacher, I would like children not only to learn about hard subjects such as ethical theories but also to learn to think critically and be able to apply it to everyday life. Recognizing the complexity of the world we live in; I want to prepare my future students for the real world. This includes thinking about and handling with challenging topics of real life, some of which might really challenge our moral standpoint and even our way of life.

The framework of Effective Altruism is a proposition to approach hard real-life questions and address the massive problems we are currently facing as humanity. It not only highlights the moral catastrophes going on in the world, from extreme poverty to factory farming and accentuates our lack of preparedness for eventual catastrophic events such as pandemics but it is also a call to action. This is why I consider Effective Altruism a valuable topic to discuss with future pupils in ethics class.

3 Defining Effective Altruism

The core idea of the Effective Altruism movement goes back to 1972 Peter Singer's *Famine, Affluence and Morality* among other influential thinkers. Singer argued in his book, that given the vast amount of suffering taking place in the world, such as through famine, we should give large proportions of our income to help those most unfortunate (Singer, 2015, p. 13). As Singer himself described, it took many years until he started to see significant efforts being made to helping the poorest people in the world. The work of Abhijit Banerjee and Esther Duflo, professors at MIT, in the field of poverty alleviation developed into Poverty Action Lab founded at the start of the century. The organization is dedicated to find the most effective (evidence-based) interventions in the fight against poverty using the scientific method. In 2006 Holden Karnofsky and Elie Hassenfeld, two hedge fund analysts in New York were faced with a problem that ultimately changed their lives. As they and six other colleagues were trying to determine where to donate to, they realized that figuring out which organizations would do the most good with their donations was not an easy question to answer. Unable to find evidence-based interventions they were willing to fund even after their plights for this type of data to the organizations themselves remained unanswered, they decided to fund GiveWell (Singer, 2015, pp. 13–16). It is a non-profit organization (NPO) dedicated to the evaluation of programs in human health that stand out the most. They regularly check with the organizations they evaluate on updates and publish their analysis of each one of them in full detail on their website (GiveWell, 2018).

In 2009 William MacAskill and Toby Ord, another philosopher from Oxford who donates everything he earns above \$20,000 a year, founded Giving What We Can. It is an international charity dedicated also to evaluate charities similarly to GiveWell and also encourages people to pledge a share of their life incomes, usually 10 percent, to effective charities. In line with students' career advice, MacAskill and five friends founded 80,000 Hours, an organization dedicated to research the impact of various careers to give EA-aligned career counseling. It is called 80,000 Hours because this is roughly the amount of hours people spend working in their careers (Singer, 2015, pp. 16–18).

The term 'Effective Altruism' was first established through the foundation of the Center for Effective Altruism in 2011; at a time where the movement was just in its beginnings and there was no official definition of the term. Today, ten years later there are multiple definitions and understandings established all around the world (MacAskill, 2019b, p. 12; Singer, 2015, p. 18).

The Center for Effective Altruisms defines it on their website as follows: “Effective altruism is about using evidence and reason to figure out how to benefit others as much as possible, and taking action on that basis“ (CEA, n.d.). The Effective Altruism community, is according to them, “a global community of people who care deeply about the world, make helping others a significant part of their lives, and use evidence and reason to figure out how best to do so“ (CEA, n.d.). William MacAskill, one of the founders of Effective Altruism explained in his book *Doing Good Better* that Effective Altruism is about asking:

How can I make the biggest difference I can?’ and using evidence and careful reasoning to try to find an answer. It consists of the honest and impartial attempt to work out what’s best for the world, and a commitment to do what’s best, whatever that turns out to be. (MacAskill, 2016, p. 11)

Robert Wiblin, an important thinker of the Effective Altruism movement who works at 80,000 Hours, sees EA as “the use of evidence and analysis to take actions that help others as much as possible” (Wiblin, 2015). For Peter Singer EA is “[...] a philosophy and social movement which applies evidence and reason to working out the most effective ways to improve the world [...] (Singer, 2015, p. 4). The CEA also gives the following definition in their introductory website into Effective Altruism:

[...] a research field which uses high-quality evidence and careful reasoning to work out how to help others as much as possible. It is also a community of people taking these answers seriously, by focusing their efforts on the most promising solutions to the world’s most pressing problems. (CEA, 2020)

All definitions have some aspects in common. They focus on the idea of maximization and define the movement’s goal as ‘doing good’, ‘helping others’ or ‘improving the world’ in a general sense. They also introduce the idea that in order to do so, the Effective Altruism community uses ‘evidence’ and ‘reason’ to find out the ‘best ways’ or the ‘most effective ways’ to help others. According to MacAskill (2019b, p. 13) the last definition from the CEA (CEA, 2020) could be understood as the most largely accepted definition of Effective Altruism, formally endorsed by the large majority of leaders in the EA community.

However, as MacAskill (2019, p. 13) also noticed, this definition's primary audience is the general public and not the philosophical academy. So, he proposed a more elaborated definition for academic discussion on the topic of Effective Altruism:

Effective Altruism is:

(i) the use of evidence and careful reasoning to work out how to maximize positive impact on others with a given unit of resources, tentatively understanding 'the good' in impartial welfarist terms, and

(ii) the use of the findings from (i) to try to improve the world.

(i) refers to effective altruism as an intellectual project (or 'research field'); (ii) refers to effective altruism as a practical project (or 'social movement')

(MacAskill, 2019b, p. 14)

He argued that Effective Altruism consisting of the intellectual and the practical projects and not on a set of normative claims, is *non-normative*. The goal of the movement is to *maximize* its members' impact by doing as much good as possible with the resources available to them. To do so the EA community makes use of scientific principles and evidence-based evaluations. And at last, MacAskill (2019, p. 14) claimed Effective Altruism to be *tentatively impartial and welfarist* – "As a tentative hypothesis or a first approximation, doing good is about promoting wellbeing, with everyone's wellbeing counting equally" (MacAskill, 2019b, p. 14).

For MacAskill Effective Altruism is to moral philosophy, as engineering is to science. While science and moral philosophy are concerned with the discovery and understanding of general truths about the world and morality, respectively, engineering and Effective Altruism are the application of science to build systems beneficial for society (MacAskill, 2019b, p. 11f).

4 The ethical argument for Effective Altruism

On your way to work, you pass a small pond. On hot days, children sometimes play in the pond, which is only about knee-deep. The weather's cool today, though, and the hour is early, so you are surprised to see a child splashing about in the pond.

As you get closer, you see that it is a very young child, just a toddler, who is flailing about, unable to stay upright or walk out of the pond. You look for the parents or babysitter, but there is no one else around. The child is unable to keep her head above the water for more than a few seconds at a time. If you don't wade in and pull her out, she seems likely to drown. Wading in is easy and safe, but you will ruin the new shoes you bought only a few days ago, and get your suit wet and muddy. By the time you hand the child over to someone responsible for her, and change your clothes, you'll be late for work. What should you do? (Singer, 2010, p. 3)

This is the thought experiment Australian philosopher Peter Singer challenged the students in his university course on Practical Ethics with, when they start to address the topic of global poverty (Singer, 2010, p. 3). The reality is that around 14,300 children as the one in the thought experiment, die every single day (World Health Organization, 2020). They surely do not drown in ponds in affluent countries, but rather die from easily preventable diseases, such as measles, malaria, and diarrhea. Conditions people in developed nations have not experienced for decades, and if they do, they most certainly do not have the same fatal results as in developing nations. These diseases arise due to bad living conditions, from lacking access to safe drinking water and or a lack of sanitation, undernourishment, but also mainly due to mere poverty, as parents cannot afford any medical treatment once their children fall ill (Singer, 2010, p. 3).

Peter Singer's argument goes as follows: if you are willing to jump into a pond to save a child's life and in the process ruin your suit and shoes, why are you not willing to commit the same sacrifice for a child dying of malaria or diarrhea? Instead of buying your next pair of shoes or a new garment, you could instead donate the money to an organization working to reduce poverty and provide basic services in health care and infrastructure, such as access to potable water (Singer, 2010, p. 5):

“[W]e all spend money on things we don't really need, whether on drinks, meals out, clothing, movies, concerts, vacations, new cars, or house renovation. Is it possible that by choosing to spend your money on such things rather than contributing to an aid agency, you are leaving a child to die, a child you could have saved?” (Singer, 2010, p. 5)

This argument challenges us to recognize that the money we spend in non-essential day-to-day expenses could actually save someone's life if donated to charity helping those most in need. Philosopher Peter Unger further developed Peter Singer's thought experiment as follows:

„Bob is close to retirement. He has invested most of his savings in a very rare and valuable old car, a Bugatti, which he has not been able to insure. [...] Not only does Bob get pleasure from driving and caring for his car, he also knows that its rising market value means that he will be able to sell it and live comfortably after retirement. One day when Bob is out for a drive, he parks the Bugatti near the end of a railway siding and goes for a walk up the track. As he does so, he sees that a runaway train, with no one aboard, is running down the railway track. Looking farther down the track, he sees the small figure of a child who appears to be absorbed in playing on the tracks. Oblivious to the run-away train, the child is in great danger. Bob can't stop the train and the child is too far away to hear his warning shout, but Bob can throw a switch that will divert the train down the siding where his Bugatti is parked. If he does so nobody will be killed but the train will [...] destroy his Bugatti. Thinking of his joy in owning the car and the financial security it represents, Bob decides not to throw the switch.“ (Singer, 2010, p. 13f)

This variation of the thought experiment adds a factor we commonly face when thinking about our contribution in the fight against real-world poverty: the uncertainty about the results of our sacrifice. If the child hears the train soon enough to get to safety Bob's eventual sacrifice would have been for nothing. Similarly, we might wonder whether our donation to charity might not

fulfill its intended purpose (Singer, 2010, p. 14). As Peter Singer (2010, p. 14) highlights, most people would almost always condemn Bob's (in)action, as we should not risk a child's life only to save a car, no matter what type of car. However, "[b]y implication, we should also believe that with the simple act of saving money for retirement, [...] we are effectively refusing to use that money to help save lives." (Singer, 2010, p. 14). This implication is certainly not easy to digest. But most people do feel an obligation to help and condemn those who do not, when prompted to think about real-life individuals suffering, even at a great cost to ourselves (Singer, 2010, p. 15).

Even in a situation where it is not directly about life and death, most people agree to aid and sacrifice something in the process, as a further example devised by Unger depicts:

“You are driving your vintage sedan down a country lane when you are stopped by a hiker who has seriously injured his leg. He asks you to take him to the nearest hospital. If you refuse, there is a good chance that he will lose his leg. On the other hand, if you agree to take him to the hospital, he is likely to bleed onto the seats, which you have recently, and expensively, restored in soft white leather.” (Singer, 2010, p. 14f)

4.1 In the classroom

These are three thought experiments easily presentable in ethics class and ideal to start a discussion about the student's thoughts and intuitions of the moral imperative to help others. It is especially important to stay impartial and not give the students the feeling that there is a “right” or “wrong” answer. This is not only relevant as such but would also respect the indoctrination prohibition, which is further discussed in chapter 8. The students are to be encouraged to think for themselves, find and define their standpoints, develop a solid argumentation for further discussion with others. But through this, pupils should also learn to recognize mistakes in their logic and learn to update their standpoints when fitting. However, it is still important for the teacher to highlight faulty or inconclusive arguments; non-indoctrination does not mean, that all standpoints and arguments put forth in class are equally “right” and “acceptable”.

These thought experiments also offer points of critic that students should attempt to find and elaborate on. The question on the validity of the conclusions of the thought experiment for the real world is one example of a point of critic. Students could also be encouraged to find and develop new thought experiments they may find more fitting. A lesson plan discussing the exact implementation of these thought experiments in class can be found in chapter 8.

4.2 The logical argument

The topic can then be further elaborated by dissecting the thought experiments into its single parts, mainly two premises: (1) Suffering and death are bad; and (2) if you are able to help to prevent (1) from happening, even at a certain cost, you ought to do it. These two premises can be then further elaborated into a logical argument, as Peter Singer (2010) did.

Seemingly, our intuitive belief is that at least in a situation where we are the only one able to save or help someone else, we ought to do it, even if it comes at a great expense for us. Peter Singer then went on to evaluate the matter beyond intuition and made the case for helping those under poverty with the following logical argumentation (Singer, 2010, p. 15):

First premise: Suffering and death from lack of food, shelter, and medical care are bad.

Second premise: If it is in your power to prevent something bad from happening, without sacrificing anything nearly as important, it is wrong not to do so.

Third premise: By donating to aid agencies, you can prevent suffering and death from lack of food, shelter, and medical care, without sacrificing anything nearly as important.

Conclusion: Therefore, if you do not donate to aid agencies, you are doing something wrong. (Singer, 2010, p. 15f)

The thought experiments elaborated above are all scenarios epitomizing this logical argument. Clearly, saving someone's life is more important than either ruining your new pair of shoes or losing an expensive car even if it is your savings for retirement. This human drive to help others seems to remain true even when it comes down to less precarious situations, such as the thought experiment of the wounded hiker shows. Most people disregard the costs of the reupholstering as insignificant compared to the cost of someone losing a leg (Singer, 2010, p. 16).

Certainly, most people would not deny the first premise, that suffering and death from lack of food, shelter and medical aid are bad. In 2019, 690 million people – a number that is expected to rise significantly due to the Covid-19 pandemic were living in extreme poverty and thus lack the means and access to even the most basic medical services (The World Bank, n.d.-c). They live in the most precarious conditions, short on food and without access to safe drinking water, unable to send their kids to school or even bring them to the doctor when sick. In fact, in 2019, around 5.2 million children under five years died from mostly preventable and treatable causes such as birth complications, pneumonia, diarrhea and malaria (World Health Organization,

2020). The seriousness of the situation is certainly hard to grasp, especially without being exposed to it on a day-to-day basis.

However, as Peter Singer's argument goes, putting yourself in the place of others, even those on the other side of the world, unable to save their children due to a lack of resources or even the children themselves, "[...] is what thinking ethically is all about." (Singer, 2010, p. 16). He explained that this principle is widely known as the *Golden Rule* - "[d]o unto others as you would have them do unto you" (Singer, 2010, p. 16). This ideal calls us to value the needs and desires of others as if it were our own, and when it comes to the poorest individuals, their desire to satisfy basic human needs as well as to save their children should count as if it were our own (Singer, 2010, p. 16f).

The second premise is also almost natural, as illustrated by the three thought experiments elaborated above. The only excuse to not do our best to prevent something terrible to happen, is if the process puts something nearly as important at risk., according to Singer. He recognized that the term *nearly as important* is surely vague, but it is hard to come up with something that is as valuable as saving a child's life (Singer, 2010, p. 17).

The two first premises together bring the thought experiments to life. Here, we have stressed out that there are not only people in great need of help, but that we also ought to help them, as long as it does not require us to make an enormous sacrifice. The third premise is just the means for fulfilling this purpose. There are enough organizations on-site helping those in need, many of them funded by donations from individuals in industrialized nations.

4.3 In the classroom

For pupils to be presented with Singer's logical argument, they need to first be competent in areas of logic and argumentation theory. Only when the students understand what premises, conclusions and logical fallacies are, and know how to use them correctly, does the discussion of Singer's logical argument make sense. As discussed above in regard to the thought experiments, students should be encouraged to develop and present their own thoughts and arguments. A lesson plan on this sub-topic is presented in chapter 8.

4.4 Among the 10 percent

Usually when people are confronted with this argument and its conclusion, a certain level of opposition arises, as it becomes clear that the money to buy any type of luxury, no matter how small, could be donated to charity instead (Singer, 2010, p. 18). Some will say that they do not have the luxury to donate money, as they are just living a simple middle-class life and that the responsibility to help lies on *the 1 percent*. This designation first went mainstream during the Occupy Wall Street movement in 2011 to refer to the top 1 percent of income earners in affluent nations, primarily in the United States (MacAskill, 2016, p. 15):

The term came from a popular statistic that the richest 1 percent of the population receives 24 percent of total income – that’s more than \$340,000 per year, twelve times the \$28,000 earned by the typical American worker. References to the 1 percent versus the 99 percent – i.e., the rest of the population – quickly became shorthand for the income gap in America (MacAskill, 2016, p. 15).

Inequality is not only a problem between nations but also within nations, as the gap between the richest households and the rest skyrockets. Between 1979 and 2007, the income of a typical household in the USA grew by less than 40 percent while that of the richest one per cent grew by 275 percent (MacAskill, 2016, p. 15). In 2018, three men, Jeff Bezos, Bill Gates and Warren Buffett, owned as much as the bottom half of American (total of \$300 billion). The Covid-19 pandemic has only increased this gap, as between January 2020 and April 2020 34 of the nation’s wealthiest billionaires saw their wealth increase by tens of millions of dollars (Collins et al., 2020). This might leave those not belonging to the richest segments of the population feeling powerless. But they forget to take a broader view worth noting.

The inequality of incomes between countries is much higher than the inequality within countries (Roser, 2013). Over the years nations around the world have been experiencing a significant increase in the sizes of their middle class. At the end of 2016, there were 3.2 billion people worldwide living in or above the global middle class, with around 140 million joining annually, especially in Asia (Kharas, 2017). It is important to note, that this number will certainly be affected by the current Covid-19 pandemic. However, many countries such as the United States of America, Japan, France, Korean Republic, among others, have more than 90 per cent of their population living in or above the global middle class with high disposable incomes (Kharas, 2017).

Graph 1 illustrates the global distribution of income in 2011. Certainly, these numbers have changed somewhat over the last decade, as developing countries, especially China, have gotten significantly richer since then. However, the highest wages are still earned by people in developed countries. Those earning professional salaries in the U.S. for example, belong to the top 5 percent of earners on a global scale and some even to the top one percent (Roser, 2013; Wiblin, 2017b).

Even the values for the median household disposable income (income after income taxes) of many rich nations remain significantly above the average global income. In Austria for example, the value was slightly above €28,000 in 2018, in France it was at €23,000 and in the United States it was at \$35,600 (OECD, 2020). So, despite not being in the top one percent of the income brackets, median citizens in most of the developed nations have enough income to contribute with to help others.




Figure 1 Global Income Distribution in 2011. Graph adopted from Wiblin (2017)

We have just outlined the main argument for helping others and why even if we do not belong the richest one per cent, we ought to do it. The next question is in which way and through which means exactly should we help. Peter Singer's main thesis is that we should support

organizations working to help those in gravest need. The question is which organization exactly we should donate to. There are thousands of organizations, each one promoting their causes. With so many options, how should people decide whom to support. Effective Altruism does not only entail the argument for why we should help but also *how* we should do it.

It is important to note here, that Effective Altruism is not limited to finding the organizations that would help those in gravest need, but goes well beyond that, as discussed above. It certainly started as such but entails much more than that today, the details of which are discussed in the chapters that follow.

4.5 In the classroom

One idea to discuss this topic would be to allow the students to estimate the mean income per year per person or household for different countries (see Figure 1). This could also include estimations for their current household income and the yearly income they expect to earn in the future. In case the teacher would like to be somewhat more provocative, they could first discuss the case with the pupils on whether the richest 10 percent should be the ones responsible for helping the poor and not necessarily the median global income earners. Especially attentive students might even recognize where this discussion is leading towards to and question the teacher's statements. In such a case should these remarks be welcomed with praise, as they not only questioned the teacher's presented position, but also spoke out against the majority of the classroom should they have agreed with the teacher. At the same time, they are questioning an otherwise for them personally beneficial assumption – not having to help the poorest with their future earnings. The second point should be considered of special relevance not only for ethics class but especially for a democratic and pluralistic society. The willingness of individuals to publicly speak out their personal opinions, is according to the spiral of silence theory from Elisabeth Noelle-Neumann, dependent on their perception of the opinion climate. The more people are perceived to share one opinion, the lower is an individual's willingness to stand publicly against it (Noelle-Neumann, 2001).

It is certainly important to clarify the presented statements by the teacher, even if students fail to remark the error in the teacher's presumed thinking. The spiral of silence theory can still be shortly discussed; even if it is to discuss the students' failure to contradict the teacher and their schoolmates. The theory as such highlighting an important ethical dilemma is incredibly valuable to be discussed in ethics class as such.

5 Five key aspects of Effective Altruism

With so many problems in the world in need of attention, Effective Altruism deals with the question where to begin. Knowing, that our monetary resources and time to help are limited, where are these resources currently most needed? What should be prioritized and how could we help? These are the leading questions of the Effective Altruism community, which has set itself the goal to determine where we can do the most good.

Effective Altruism has adopted a common framework aligned with common scientific practices to determine how to prioritize between causes. To determine how much impact work in a given field would have, they consider its scale, neglectedness and solvability. The fields in which interventions would make the biggest difference are: great in scale, highly neglected, and highly solvable or tractable (EffectiveAltruism.org, 2020). MacAskill elaborated these three and two other aspects used by the Effective Altruism community to decide which cause area to focus on in his book *Doing Good Better* (2016). The next chapters entail a detailed analysis of all five aspects listed below, and are mainly guided by MacAskill's (2016) work:

- Scale: How many individuals benefit, and by how much?
- Effectivity: Is this the most effective thing you can do?
- Neglectedness: Does this area already receive enough attention?
- Alternative: What would have happened otherwise?
- Expected value: What are the chances of success, and how good would success be?

5.1 Scale

How many individuals benefit, and by how much?

During situations of crisis, such as war, natural catastrophes or a pandemic, when medical resources become scarce and there is no time to lose, medicals make use of the principle of *triage*. This is a system that helps medics determine where to set their focus, i.e., prioritize between so many casualties. Patients are classified according to the degree of urgency of their situation so that medics can first treat those in most critical circumstances. This also prevents them from putting their efforts into someone whose situation is irretrievable. This might sound harsh and even cold to some, but with so many casualties and only a limited set of resources, medics know they cannot save everyone. To save the most people they can, they need to prioritize and that entails making tough choices. An example of such a situation demanding choices between life and death was the Rwanda genocide in 1994, during which 800,000 men,

women and children were exterminated at the hands of fellow citizens. Medics such as James Orbinski were forced to act fast and make hard triage choices to save as many people as possible (MacAskill, 2016, p. 32).

5.1.1 A constant situation of crisis

Taking a look at the vast numbers of people around the globe dying from easily preventable deaths, the EA community argues that we find ourselves in a crisis of similar magnitude every single day (MacAskill, 2016, p. 32). The mere fact that we are not directly confronted by those suffering, as Orbinski was in 1994, does not make the situation less real for the individuals affected. And according to Peter Singer's argument evaluated in chapter 4, it also does not make our responsibility nor our ability to help less real either.

However, we have only limited resources to give for helping those in need. Every single individual can only give so much time and so much money to helping others. So, "[o]f all the ways in which we could make the world a better place, which will do the *most* good? Which problems should we tackle immediately, and which should we leave for another time?" (MacAskill 2016, p.32). It is of utmost importance to evaluate all the possible actions and compare them to one another to identify how many people benefit from the help, and by how much (MacAskill 2016, p.32).

5.1.2 Choosing which charity to support

Answering this question requires the understanding of the consequences of each possible action. To illustrate this matter, MacAskill (2016, p. 32) discussed which of all the charities there are should receive your donation. To decide you need to know what exactly the charities plan to do with your donation. The problem is that many non-profit organizations (NPOs) do not give their donors clear answers on how much impact their donations are having or even where exactly they are invested in. Many have no information on how much their programs cost and thus where your donation would be likely to be invested in. This might not surprise many of us, as we are used to this type of secrecy, but imagining this same type of uncertainty in any other scenario, where money is involved, is unthinkable (MacAskill, 2016, p. 32 f).

There are, indeed, some organizations that provide some information on what exactly your donation is financing, such as a specific number of books they provide for people in need. However, this does not tell us anything about the impact these items are having in the lives of

those who receive them (MacAskill, 2016, p. 34). What benefit do they have from getting those books in the first place? Does it help them understand the world better? If we are talking about schoolbooks, aimed at helping children do better in school, is this really the best way to do so? How much impact are they having?

These are the questions that shift our thinking from “what and how many items/activities” am I sponsoring to “how are those items/activities” I sponsored are “actually improving lives?” Only this way are comparisons between interventions really possible, by measuring “[...] impact in terms of the size of the benefits we confer through those actions” (MacAskill, 2016, p. 34).

5.1.3 A situation of crisis requires us to set priorities

There are clearly cases, where these comparisons are easier to make. For Orbinski’s critical situation in Rwanda, it was clear that saving someone’s life was more important than saving someone from losing a leg. It was also clear for him, that the benefit of saving five lives by providing simple lifesaving surgeries was bigger than saving one life only savable through a much more complex time-consuming surgery (MacAskill, 2016, p. 34).

However, most of us do not find ourselves directly confronted with such critical situations. We even have the option to decide whom to help of all the people in need worldwide. This, clearly, makes the evaluation much more complicated. How are we going to compare helping fund cancer research to saving someone from blindness? Should we focus on finding a cure for Alzheimer or donate to the fight against malaria?

5.1.4 Assessing the impact of health interventions: the QALY

When it comes to health, researchers have dedicated years to answer questions like these. And for over forty years now, they have been using a metric called the QALY, short for quality-adjusted life year, which helps make comparisons among different health interventions (Sassi, 2006). The QALY reflects the two ways you can give a health benefit to someone. You can save someone’s life, ergo, extending it, and you can improve someone’s life during their lifetime (MacAskill, 2016, p. 34). These two types of benefit are combined in the QALY, “[...] using survey data about the trade-offs people are willing to make in order to assess how bad different sorts of illnesses or disabilities are” (MacAskill, 2016, p. 35).

As an example:

“[...] people on average rate a life with untreated AIDS as 50 percent as good as life at full health; people on average rate life after a stroke as 75 percent as good as life at full health; and people on average rate life with moderate depression as only 30 percent as good as life in full health.” (MacAskill, 2016, p. 35)

The QALY can be illustrated in graphs, such as the one below. Here, “L is the residual life expectancy of the individual at age a, and the represents individual years within that life expectancy range” and “[...] Q is the health-related quality of life weight attached to the relevant year of life” (Sassi, 2006, p. 403). There, the life expectancy and quality of life of a patient without treatment would be $(a+L, Q)$. Were the specific health intervention (i) to take place, these values would increase to $(a+L^i, Q^i)$ respectively (Sassi, 2006, p. 403).




Figure 2 Health profiles both with intervention i (solid line) and without intervention (broken line). Figure adopted from Sassi (2006, p. 404).

The QALY helps us quantify the benefits a health intervention in both years of life expectancy and or quality of life can bring to someone (Sassi, 2006). MacAskill (2016, p.36) provided a very simplified calculation to help the reader grasp the concept. E.g., improving someone’s life by 20 percent for the next 60 years of their life results in twelve QALYs ($60 \times 20\% = 12$); extending someone’s life living currently at 70 percent health, due to any type of circumstances, by ten years results in seven QALYs ($10 \times 70\% = 7$) (MacAskill, 2016, p. 36).

There are several other metrics used to evaluate health interventions, besides the QALY, e.g., the DALY, for disability-adjusted life years (Sassi, 2006), and the WALY, calculating in well-being adjusted life years (MacAskill, 2016, p. 39). These metrics are commonly used by health economists to assess the cost-effectiveness of different health interventions. They evaluate how much it costs to fund different programs of health intervention and convert the results into QALYs. This way, they are able compare different health treatments and recognize which

programs bring the largest benefits for a given amount of money. Given that we have limited resources and we have argued why we should try to make the best out of our resources, then, *ceteris paribus*, we should spend those resources in the program that brings the most QALYs (MacAskill, 2016, p. 37).

5.1.5 Limitations of the QALY

Clearly, these metrics are not the ultimate solution to this complex problem. Based on survey data, it is prone to human biases and misestimations on how much their life would be affected by a given health deficit. We know that patients themselves tend to evaluate their conditions as less bad than the general public does. However, it is hard if not impossible to know the exact origin of this miscalculation. Is it based on an overestimation by the broad public on the negative impact a given medical condition would be? Or do already affected patients subconsciously set their 100 percent health mark lower than they would have done without the condition? (MacAskill, 2016, p. 38).

The difficulty in answering these types of questions keep academics busy. There are other aspects to consider as well, such as whether it is especially important to prevent the deaths of younger individuals or give people with difficult backgrounds priority. These issues will certainly not be solved in the near future (MacAskill, 2016, p. 38).

5.1.6 The impact gap between health interventions

However, as MacAskill (2016, p. 38) argued, exact numbers on how bad or good a given condition is, tend not to be as important as one might think. They at least help to make important comparisons between programs, at least in principle. As a matter of fact,

[...] programs differ dramatically in how great an impact they have, so even a rough idea of how many people are affected, and by how much, is often enough to show that one program has a much larger impact than another. (MacAskill, 2016, p. 38)

There are, indeed, some organizations that have, according to the QALY (and other health intervention metrics), multiple times more impact than others. The differences lie in the fields of action of the programs themselves. The most cost-effective organizations treat health conditions with terrible effects on human lives for a very low cost. It is not a coincidence that these organizations tend to work in the poorest parts of the world. People there often lack basic, inexpensive things that once provided make an immense difference. They also suffer from

easily treatable and or preventable diseases as further elaborated in chapter 6. So, it is just understandable that programs treating these types of conditions are several orders of magnitude more cost-effective at improving lives than many local programs (GiveWell, n.d.-c).

For example, let us say, that you want to donate \$50 to charity. A friend of yours recommends you the Make A Wish foundation, which grants children with terminal diseases a wish, such as meeting their role model or going to Disneyland. According to their website, the average cost of granting a wish for one child in the U.S. in 2018 was \$11,161 (Make A Wish, 2018). However, after asking your sister, she tells you she recently donated to Evidence Action's Deworm the World Initiative, which administrates deworming medication to children in developing countries. According to GiveWell's calculations, a *meta-charity*, dedicated to evaluating nonprofit organizations, it costs about \$0.66 (\$0.46 excluding government contributions) to deworm a child in Kenya for one year (GiveWell, 2020b).

So, you can either donate those \$50 to Make A Wish Foundation and contribute to 0.4 percent of the costs for the fulfillment of a child's wish in the U.S. Or you could donate the same \$50 to Evidence Action's deworming initiative and fund the deworming of 108 children in Kenya. The contrast between both charities in what your donation can achieve is just undeniable.

5.1.7 Making the right questions

In summary, this chapter covered two aspects on the question of *scale*. First, the need to consider the consequences of our actions: How are my actions improving peoples' lives? And second: How many individuals am I actually helping and by how much? These questions are crucial not only when considering where to donate to, but also when it comes to volunteering, deciding to buy ethical produce, and even in the arduous decisions when choosing a career (MacAskill, 2016, p. 42). The next aspect of Effective Altruism focuses on the importance of finding the best activities to help others.

5.2 Effectivity

Is this the most effective thing you can do?

Over the years there have been multiple critics of development aid. Dambisa Moyo, author of *Dead Aid* (Moyo, 2010) and William Easterly with *The White Man's Burden* (Easterly, 2007) are just a few examples. They argued that despite the West giving out trillions of dollars for development aid, there has not been much progress in the receiving countries (Moyo, 2010). Their main criticism is that aid allegedly hinders people from finding their own solutions while feeding corrupt politicians and undermining local institutions. They see the establishment of free markets as the only solution to bring people out of poverty (Banerjee & Duflo, 2012, p. 3f).

5.2.1 PlayPump: an aid program gone wrong

Of course, there is a track record of aid programs gone wrong. A good example for this is PlayPump International, set up by Trevor Field and his father-in-law, two South African men. Their urge to help started during a fishing trip, where they watched women of a rural village waiting for hours in front of a windmill-powered water pump for wind to come. Shocked by that image, Field started a project to set up a special playground merry-go-round for children, whose generated power would be utilized to pump underground water to a storage tank people could access. A World Bank Development Marketplace Award in 2000, attracted funding and international attention, including a collaboration with the Caser foundation. The program was praised by politicians, like Bill Clinton and celebrities, such as Jay-Z raised thousands of dollars for the cause. By 2009, PlayPump International had set up 1,800 PlayPumps across South Africa, Mozambique, Swaziland, and Zambia (MacAskill, 2016, pp. 1–3).

However, shortly after, two reports were released, one by UNICEF and one by the Swiss Resource Centre and Consultancies for Development (SKAT), both outlining how unpractical the pumps turned out to be. Unlike the typical playground merry-go-rounds, which would spin freely once they gain momentum, the PlayPumps need constant force to rotate. Children not only got exhausted and ended up hurt, but in some cases, they were even paid to use the pump. At the end, women were the ones that ended up pushing the merry-go-round themselves, which they found not only tiring but demeaning. The pumps also broke down shortly after installation and the community did not obtain the technical support they were promised (MacAskill, 2016, p. 4).

Even more troubling was the following finding:

What's more, no one had asked the local communities if they wanted a PlayPump in the first place. When the investigators from SKAT asked the community what they thought about the new PlayPump, many said they preferred the hand pumps that were previously installed. With less effort, a Zimbabwe Bush hand pump of the same cylinder size as a PlayPump provided thirteen hundred liters of water per hour—five times the amount of the PlayPump. A woman in Mozambique said, “From five A.M., we are in the fields, working for six hours. Then we come to this pump and have to turn it. From this, your arms start to hurt. The old hand pump was much easier.” One reporter estimated that, in order to provide a typical village's water needs, the merry-go-round would have to spin for twenty-seven hours per day. (MacAskill, 2016, p. 4)

The PlayPump was not only inferior in almost all relevant aspects to the traditional hand pumps, but it was also four times more expensive. Soon after the reports came out, the world turned their backs from the PlayPumps and even the Case Foundation publicly acknowledged the failure of the program and shut down its U.S. base. Despite this, Field's nonprofit lives on, under the name of Roundabout Water Solutions and continues installing the same model of PlayPumps across South Africa (MacAskill, 2016, p. 5). This is an extreme example of development aid gone wrong. The project was not only ineffective in addressing the problem, but it actually did more harm than good.

Now, let us take a step back and see the bigger picture. As Banerjee and Duflo (2012 p.4) depicted, the debate about whether development aid in general is *good* or *bad* is not the right question to make as the subject is too complex to give the easy answers this question imposes. It is more about questioning what types of projects the aid is funding. No one would disagree that we should help poor people when we can. However, as we saw in the PlayPump case, understanding exactly *what* we are supporting is the key question. Are the projects in place actually effective in helping those who need it?

5.2.2 Aid programs changing the world for the better

A good example for an effective aid program was the eradication of smallpox in the 20th century. It was a horrible very infectious disease, which affected all levels of society and ended up deadly for up to 60 percent of those who got infected. Survivors were left badly scared from the disease at best and in some cases even blind. It was particularly deadly for children, killing up to 98 percent of infected children in Berlin during the 19th century (Riedel, 2005, p. 21). In the 1950s, the vaccination program against smallpox eradicated the disease successfully in many areas of Europe and North America. In 1967, a global campaign was set in motion by the World Health Organization, leading to the ultimate eradication of smallpox in the rest of world by 1977 (Riedel, 2005, p. 25). Today, it is unimaginable that smallpox killed around three hundred million people in the 20th century alone (Henderson, 2011).

5.2.3 Finding the needle in the haystack with the help of science

Now, how exactly are we going to determine whether a health intervention will end up having extraordinary consequences, such as the eradication of smallpox, or at worst, causing more trouble to the poor as the PlayPumps did? For this, we need to evaluate every single program and compare them to one another to spot the ones actually making the biggest difference.

A good method for doing this, is through randomized control trials (RCTs), a standard experimental setting in science. This setting randomly assigns test subjects either to one or more experimental groups or a control group, allowing to compare the interventions afterwards and determine which ones actually work and by how much. This is a common process in all areas of medicine, especially in the development of new drugs and treatments as a standard procedure in the pharmaceutical industry (MacAskill, 2016, p. 7). However, it took many years and intersectional collaborations to come to the idea to apply RCTs in development economics to test the effectiveness of various interventions.

5.2.4 Improving school attendance in Kenya

Michael Kremer, an American development economist, went to Kenya in the 1990s to visit a family he had lived with for a year while teaching in English at a secondary school. There he met with Paul Lipeyah, a friend working for the Dutch charity International Christian Support, today known as Investing in Children and Theis Societies (ICS). ICS's main program attempted to improve school attendance and test scores by providing a package of different school specific supplies, such as free textbooks, school uniforms and even additional teachers. Before rolling out the program to seven new schools, Kremer urged Lipeyah to test the program by using randomized control trials (MacAskill, 2016, p. 7).

With support of collaborators, Michael Kremer and Edward Miguel, an economist, set up the RCTs to test the program, providing only some of the schools with additional textbooks. The results were a bummer, as only high-achieving students seemed to have profited from the additional textbooks. In a second trial they ran the same experiment with flipcharts, allowing teachers to adapt the classes to their students' needs, without effect. They then tested increasing the number of teachers. Again, without success. In further experiments, they tested other intuitive interventions to improve the students' situation, but all attempts remained fruitless. However, Kremer did not give up in his attempt to find an intervention that actually worked (MacAskill, 2016, p. 7f).

5.2.5 The best way to help might not be the most intuitive idea

One day, a friend of Kremer at the World Bank suggested them to test deworming the children. Most people in developed nations do not know about intestinal worms, parasites that burden more than one billion people worldwide. Their symptoms are not as aggressive as malaria or AIDS, and they do not kill many people (at least not directly). But they hinder their hosts from absorbing nutrients from the food they consume, reducing their cognitive and physical development (WHO, 2016) and they were extremely prevalent in the early 2000s as Kremer was running the experiments (Evidence Action, 2020). The good news is that deworming drugs, developed in the 1950s, are extremely cheap and can be easily distributed at schools. They are also long-lasting as the effects of one single unit lasts for about a year, making it even more practical to use (MacAskill, 2016, p. 8).

5.2.6 Not only effective but extremely effective

Published in 2003, the results of the study testing the deworming treatment were unexpectedly good. “It turned out to be one of the most cost-effective ways of increasing school participation” (Kremer, as cited by MacAskill, 2016, p.8). In comparison to the control group, the treatment not only improved the children’s health and education outcomes, but also reduced absenteeism by 25 percent, one of the biggest problems in schools in Kenya. There were also spillover effects, as the treatment of whole schools reduced the prevalence of the parasitic worms in the children’s communities (Evidence Action, 2020; Miguel & Kremer, 2003). They also decreased the risks of illnesses intestinal worms usually cause such as anemia, intestinal obstruction and a weakened immune system, which increases the children’s risk of catching other diseases, like malaria (MacAskill, 2016, p. 9).

These results brought a new energy in the field of development economics. Randomized control trials became mainstream, with multiple development programs being put up to the test across hundreds of trials. Backed up by international support from organizations and governments, Michael Kremer and his wife, Rachel Glenneerster cofounded the nonprofit Deworm the World Initiative in 2009. Today, the organization provides technical assistance to “[...] governments in Africa and Asia to provide deworming treatment to over 280 million children annually, and in total have supported provision of over a billion deworming treatments since 2014” (Evidence Action, 2021).

5.2.7 Treatment with long-lasting positive effects

To understand the long-term effects of the treatment, Kremer, Miguel and colleagues followed up with the children every five years over a time span of 20 years. This is one of the first longitudinal studies of this kind and the results were published in August 2020 (Evidence Action, 2020; Hamory et al., 2020):

The authors found that children who received deworming treatment for two or three years in a row had better quality of life decades after the fact. When compared to those who had fewer years of treatment, they reported 14% higher household consumption rates and 13% higher income. In lower income countries like Kenya, these types of increases tend to represent improvements to the individual’s and their family’s diet, access to education and health, housing, and overall living conditions (Evidence Action, 2020).

Since their first study published in 2003, highlighting the effectiveness of deworming to improve school attendance, many more studies have taken place. One of them, taking place in Madagascar, found one intervention to have even bigger effects in school attendance than deworming: informing parents on the expected returns of education, either via a role model or statistics (Nguyen, 2008). The comparison is depicted in Figure 3.




Figure 3. A comparison of the size of the effects in additional school years per US\$ 100 invested in different programs. The results are based on the evaluation of different treatments by different studies. Figure adopted from Figure 6 from The World Bank (2012, p. 18).

All programs listed in Figure 3 are aimed to increase children's school attendance in developing nations. But there are two that work by several orders of magnitude better than the rest. The same \$100 could either result in 40 additional years of school attendance or only .02 years, depending only in the form of intervention obtaining the additional funding. This means that the two best programs, informing parents on the expected return and deworming children at school are 2,000 and 1,430 times more effective than conditional cash transfers (CCTs) for primary education in Mexico, respectively. So, when it comes to helping others, there is a big difference between a good use and a great use of money.

5.2.8 A fat-tailed distribution in effectiveness

For making comparisons across fields, researchers make use of metrics such as the QALY described above. This allows them to make estimates on the cost-effectiveness between health treatments. Clearly, these estimates do not provide the ultimate answer to this complex problem. However, Figure 3 is a prime example for most comparisons between programs in the foreign aid sector. In this case all interventions shared the same goal: increasing school attendance of students in poor countries, but even when making comparisons between programs with different goals, their impact seems to end up having this type of fat-tailed distributions. There seems to always be a few numbers of programs having many times more impact than the rest (MacAskill, 2016, p. 53).

So, this chapter has covered the importance of carefully evaluating the impact of different programs. Making comparisons between them allows us to determine which interventions actually make the biggest difference. It is not only about finding a program that works but finding the one that works best. The next chapter provides us with a rule-of-thumb MacAskill (2016) proposed to find the most effective programs.

5.3 Neglectedness

Does this area already receive enough attention?

With so many problems in the world, it is hard to keep an overview. Living in the West, you are more likely to see ads promoting donations to fund cancer research, than those trying to raise funds for fighting malaria. This might seem intuitive to you, but it is those problems that are taking place around us, which are the most likely to get our attention and thus our support. However, these problems are not necessarily the most pressing ones at a global scale. In fact, it is very likely that those less-loud causes acting outside our range of perception are addressing much more pressing issues than the local organizations you constantly hear about. And in part due to the lack of publicity and the lower influx of donations, they might be more resource constrained, so any donation would certainly make a much bigger difference than if you were to donate it to a better-known cause.

5.3.1 Diminishing marginal utility: the more you have it, the less it is worth

In economics, this thought is known as *marginal utility*, evaluating the value of an additional unit of a given thing. As the amount of a given resource increases, the value this resource brings you increases as well. But this relationship is not necessarily linear. In fact, the value each additional item brings you decreases the more items of the same type you already have. This is known as the law of *diminishing marginal utility*. Similar to this concept, is the *law of diminishing returns*, which describes that an increase in the amount of one single factor of production, while all other factors of production are held equal, diminishes that item's marginal return. If we take antibiotics as an example, a small dose of antibiotics might not be enough to fight an infection, so each additional unit of antibiotics will increase their effect. However, this increase in utility of additional antibiotics will only increase up to a point, its maximum total utility. Additional antibiotics beyond that point will only cause more harm than good to the patient, thus diminishing returns.

5.3.2 Increasing the value of your donation by supporting the neglected causes

So far, we have talked about the massive differences in effectiveness between programs. The best of them not only address issues that affect a high number of individuals, but they do so at a very low cost. Now, considering the law of diminishing returns, Effective Altruism encourages us to look for exactly those cost-effective programs addressing problems, that receive the least attention and thus the least resources. If a given cause already has a lot of attention and funding, then our investment in that area might not be as cost-effective as in other more neglected areas. It is exactly those projects and organizations working on those less-known causes, which are the ones we should want to support, according to EA, to help them overcome their constraints. Donating to those programs will certainly have a higher marginal utility (MacAskill, 2016, p. 58).

5.3.3 Emotions are not always the best guide

MacAskill (2016, p.58) gave a good example to depict exactly this issue. On 11 March 2011, the Tōhoku region of Japan was hit by one of the strongest earthquakes since 1900. The tsunamis reached heights of 40 meters and wandered as far as 10 km into the mainland. On 12 January 2010, an earthquake hit Haiti near Léogâne, 26 km from the country's capital, Port-au-Prince (MacAskill, 2016, p. 58; World Vision, 2019). International media coverage for both catastrophes was similarly large, dominating the news for several weeks. Everyone was reaching out to their wallets and international aid organizations offered their support. "In each case, the total international aid raised in the immediate aftermath came to about \$5 billion" (MacAskill, 2016, p. 58).

Clearly, both incidents were similar in nature, as both were caused by a massive earthquake. However, there is a key difference listed above: The scale of human suffering they caused differed dramatically. The earthquake in Japan killed almost 16,000 people, while the one in Haiti killed 250,000 people. Haiti's earthquake also injured 300,000 people and displaced at least 5 million more (World Vision, 2019).

Japan was at that point the fourth richest country in the world (today third), while Haiti was and still is among the poorest nations worldwide (MacAskill, 2016, p. 59; World Population Review, 2021). Japan was so well off, that only four days after the incident, the Japanese Red Cross published a statement saying they no longer require external assistance, either through funding or donations (MacAskill, 2016, p. 59).

There are three aspects worth noting in this example. The first aspect is our *scope insensitivity*, the term for our inability to assimilating big numbers. “When asked how much money they would donate to save 2,000, 20,000, or 200,000 migrating birds from drowning in oil ponds, people indicated they would give only slightly more to save 200,000 birds than to save 2,000 birds” (Desvousges et al., 1993 as cited by Burson et al., 2011, p. 541). We would expect our willingness to contribute to a cause to increase with the number of victims. However, that is not what happens. Our scope insensitivity is one reason for this (Burson et al., 2011).

5.3.4 The effect of an individual’s story

Second, we know that people show a so-called *singularity effect*. As an experiment at the University of Pennsylvania illustrates, humans are easily overwhelmed by the magnitude of big problems. After receiving \$5 for filling out a survey, students were asked to make a donation to Save the Children, an international charity supporting poor children. The researchers gave one group of students a flyer with the statistics on the magnitude of the problem while the other group obtained a flyer with the tragic story of one single young girl. Students gave in average \$1.16 for the first flyer and \$2.38 for the second. The authors concluded that the students were more willing to help the one identifiable child, because they felt discouraged by the magnitude of the global problem depicted in the first flyer (Banerjee & Duflo, 2012, p. 2).

The example of the earthquakes in Japan and Haiti described above showed exactly that. People were willing to help as the issue was televised and images of individual people suffering were in the daily news. However, the response was the same for both earthquakes, even as the one in Haiti was many times bigger, by one order of magnitude, to be exact.

5.3.5 Out of sight, out of mind

The third aspect was well summarized by MacAskill (2016, p. 59): “funding seems to be allocated in proportion with how evocative and widely publicized the disaster is, rather than on the basis of its scale and severity.” Around the same time, in 2008, there was a terrible earthquake in Sichuan, China. It killed 87,000 people, five times as many as the earthquake in Japan in 2011. However, they only obtained \$500 million in international aid – one-tenth that of Haiti or Japan. The earthquake in China did not receive as much media coverage as those in Haiti and Japan, which massively affected the support it got after the disaster. Having donated

to this cause would certainly have made a much bigger difference than donating to the Haitian or Japanese cause (MacAskill, 2016, p. 60).

5.3.6 We are already living in a constant situation of crisis

MacAskill (2016, p. 60) went as far as to argue that the law of diminishing returns also highlights why “in general”, donating to the best charities helping the poor makes much more sense than donating to disaster relief. Every day, thousands of people in the poorest regions of the world die of easily preventable diseases like malaria, diarrhea, tuberculosis and AIDS. This is by far a much bigger tragedy than the earthquakes described above. More than 14,000 children die every single day due to poverty, more than the number of people who died in Tōhoku (MacAskill, 2016, p. 60; World Health Organization, 2020). As MacAskill (2016) put it in *Doing Good Better*:

For every death the Japanese earthquake caused, aid organizations received \$330,000 in donations. In contrast, for every person who dies from poverty-related causes worldwide, only \$15,000 on average is spent in foreign aid and philanthropy. Partly for this reason, experts from the World Health Organization and World Bank concluded that “emergency health interventions are more costly and less effective than time-tested health activities” (MacAskill, 2016, p. 60).

5.3.7 Sudden tragic events harness our emotions

Our willingness to help those struck by natural disasters we see in the media is a clear example of our emotion-driven behavior when it comes to charity. We follow our “guts” when confronted with the event, as “[...] the emotional centers of our brain flare up: we think – *emergency!*” (MacAskill, 2016, p. 60). The thing is that there actually is an ongoing emergency happening every single day. We have only get used to it. Disasters evoke, as new events, deeper and more urgent emotions than the every-day poverty we know about but have learnt to ignore. We are thus much more willing to intervene and help those struck by this sudden tragedy than to help those suffering from poverty (MacAskill, 2016, p. 60).

So, MacAskill (2016, p. 60f) argued that a good rule of thumb in accordance with the law of diminishing returns, would be best to resist your inclination to help should you feel a strong emotional reaction to a new tragedy. There will be probably many more like you, having a similar response to the story, so your donation to that cause will probably not make as much of

a difference. Instead, he argued, you should direct your inclination to help to causes that help reduce poverty-caused misery (MacAskill, 2016, p. 60f).

5.3.8 Your donation is many times more worth overseas

“Diminishing returns also provides a powerful argument for focusing your altruistic efforts on people in poor countries rather than those in rich countries” (MacAskill, 2016, p. 61). For this he compared the costs to help a blind person in the rich world and in poorer regions. The estimates for training one guiding dog for a blind person in the U.S. are at \$50,000. When compared to health interventions in poor countries, those same \$50,000 could cure 500 people of blindness. A simple surgery can reverse the symptoms of the bacterial infection trachoma, which causes the eyelashes to damage the cornea by turning the eyelids inwards (Berkey, 2018, p. 148; MacAskill, 2016, p. 61). And as he noted:

Any program that costs one hundred dollars to cure blindness would have been fully funded in rich countries decades ago. The same is not true in poor countries, which means we can do so much more to help those people than we can at home. (MacAskill, 2016, p. 61).

5.3.9 Malaria: once a global problem, today the poor’s burden

There is a similar case when it comes to the question of how much funding certain health programs get in relation to the size of the problem they are. If we take the funding for cancer treatment and compare it to the funding malaria gets, the disparity becomes crystal clear (MacAskill, 2016, p. 61). Malaria affects around 229 million people worldwide and in 2019 alone caused the premature death of 409,000 people, mainly children (WHO, 2020). Cancer killed in 2020 alone almost 10 million people and there were 19.3 million new cancer cases (Sung et al., 2021). So, clearly cancer kills many more people than malaria, but malaria exceeds the cases by one order of magnitude (MacAskill, 2016, p. 61).

In QALY, cancer is responsible for -7.6 percent of deaths and ill health worldwide, while malaria is estimated to cost about half as many, at 3.3 percent of QALYs lost worldwide (MacAskill, 2016, p. 61). Were we to treat problems equivalent to the problem they cause, we should expect medical spending of malaria to be about half of the \$217 billion spent in cancer treatment each year. However, in reality, the annual funding for treating malaria is of \$1.6 billion – sixty times less than expected (MacAskill, 2016, p. 61).

Malaria, similar to measles is a historic disease and affected most parts of the world throughout human history. It was the Ancient Romans, who gave the disease its name *mal aria* (Medieval Italian for ‘bad air’). And as European settlers arrived in the Americas, they took the disease with them (de Castro & Singer, 2005). In fact, it was only after the United States set up their National Malaria Eradication Program in 1947, that the disease was finally terminated in the continent in 1951. Similar to the U.S., through the help of insecticide spraying, drug therapy and environmental engineering, malaria was also eradicated in Europe in the 1970s (Piperaki & Daikos, 2016). However, lacking the resources necessary to coordinate such an effort, people in poorer countries still struggle with the disease. MacAskill (2016 p. 62) recognizes here a good opportunity to focus on, as funding programs fighting malaria in poorer nations is much more strongly neglected than funding cancer treatments in the developed world.

5.3.10 Recognizing those low-hanging fruits

All examples discussed above illustrate the value of recognizing so-called *low-hanging fruits* – easily achievable wins – in our efforts to help others. This aspect of Effective Altruism attempts to help us understand that our intuition is not always rational, an aspect we have come to understand and fight in many aspects of life but not when it comes to charity. Donating to those causes everyone else does, will not have as much of an effect, as it would do if donated to a more neglected cause. This is due to the law of diminishing marginal utility: the value of each additional item brings you decreases the more items of the same type you already have. So, those resources would bring a bigger value to under-funded programs, such as the fight against malaria in poor countries.

5.4 Alternative

What would have happened otherwise?

Related to the third concept, the fourth concept of Effective Altruism might also come as somewhat counterintuitive. It is about the importance of assessing the counterfactual of a certain event: What would have happened otherwise? What are the alternate scenarios of a given event? Wanting to do as much good as possible it is important to recognize that: “What matters is not *who* does good but whether good is done; and the measure of how much good you achieve is the difference between what happens as a result of your actions and what would have happened anyway” (MacAskill, 2016, p. 70f).

5.4.1 Someone else might do a better job

To illustrate what exactly he meant, MacAskill (2016, p. 71) gave the following example: Person A collapses on the ground due to a heart attack and person B runs up towards them to perform CPR. Assuming person B to have never performed CPR before, they end up causing some damage due to their inexperience. Nevertheless, person B saves person A’s life, which is clearly a good deed. However, when running towards them, person B pushed a paramedic out of the way, who would otherwise have helped person A. The paramedic would have surely done a much better job at performing CPR than person B. In this case, it would have been better for person B to step back and let the paramedic save person A. This exemplifies that the good someone does is not a matter of direct benefit that someone causes. Instead, it is about the difference that person does. So, in this example, person B ends up causing more harm than good by hindering someone more experienced to act (MacAskill, 2016, p. 71).

5.4.2 PlayPump: causing more harm than good

As in the example of PlayPump International discussed in chapter 5.2, many programs can actually end up doing more harm than good. Trevor Field certainly had the right intentions, wanting to help poor communities easily access clean underground water and that while providing the children with new entertainment. However, the problem was that they did not consider evaluating their idea before implementing it across the country. Only later, through the assessment of independent organizations, did the harmful effects of the PlayPump come to light.

They not only installed hundreds of PlayPumps around South Africa and, Mozambique, Swaziland, and Zambia, but they even removed perfectly working conventional pumps and replaced them with PlayPumps (MacAskill, 2016, pp. 1–5). The new pumps were not only inefficient, fragile and hard to use in comparison to the conventional water pumps, but they were also four times as expensive. So, it is safe to say that the whole PlayPump program ended up causing more harm than good.

5.4.3 What are the alternatives?

This example highlights the importance of questioning the effects of a certain program as well as its alternatives. It is about taking a step back and questioning whether a certain action is really the best option there is to address a certain problem, evaluating how good an action is, but also how good it is in comparison to the alternatives (EffectiveAltruism.org, n.d.). Similar to the third key aspect of Effective Altruism, this concept also highlights the fact that putting your efforts in the most popular causes, will not make as much of a difference than focusing in other more neglected areas. The less people give a certain neglected cause attention, the bigger the difference of each single contribution will be, as the likelihood decreases that someone else might have contributed instead (EffectiveAltruism.org, n.d.).

5.4.4 Replaceability

It is exactly in this fifth chapter of MacAskill's (2016, pp. 75-78) *Doing Good Better* that he covered an important and often criticized topic in Effective Altruism: Earning to Give – a concept is also known as professional philanthropy in the Effective Altruism community (Todd et al., 2012, p. 5) Bringing up the fact that becoming a doctor is among the most popular career paths in rich countries, the contribution a single doctor does is not as high as you might expect. In this specific example: “It's the difference you make by becoming a doctor as compared to the difference someone else would make if he or she took your place” (MacAskill, 2016, p. 75).

As there are so many wanting to become doctors and only a predetermined number of places, the impact of a single person becoming a doctor is easily replaceable. MacAskill (2016, p. 75) argued that according to calculations, one additional doctor in the U.S. might actually end up saving one or two lives that would not have been saved otherwise. This is certainly a much lower number of saved lives as one might have expected.

The conclusions of this example can be extrapolated to other altruistic actions. How likely is it that someone else might be doing your volunteering activity at the local soup kitchen instead of you? Are there others wanting the same position at that non-profit you want to apply and if so, do you really think you would do a much better job? Considering these questions, help us understand that there are positions and activities that are so highly wanted that any single contribution is easily replaceable. So, our contribution there will not have much of an impact if at all.

5.4.5 Professional philanthropy: an alternative for everyone

Positions in the non-profit sector are highly contested due to their limited number. Most humanitarian causes have more than enough applicant for their job positions, but primarily lack the monetary resources to hire more people. So, what if instead of taking a job someone else might have done as well, we would make a second position for a third person possible (MacAskill, 2016, p. 75f)?

Earning to give proposes us to take up a high-earning job, from which we could donate a big portion of our salary to charity, thus enabling them to hire more people. Rather than searching for high-impact jobs in the NPO-sector, through a high-earning job, you would be able to donate amounts that would allow organizations to fill up position they would otherwise not be able to fill (Todd et al., 2012, p. 6).

There will be certainly many people trying to get those highly paid positions you would also be trying to get, and if you do not take the job, someone else certainly will. However, the fact that you would donate a big part of that salary is the key difference. Average workers donate very little, and they would most probably not try to optimize their giving decisions either. So, you would be making a huge difference by doing something that would not have happened otherwise. This way, so the argument goes, you would be doing even more good, than by working directly at the non-governmental organization (NGO) you would be donating to (MacAskill, 2016, p. 77; Todd et al., 2012, p. 6):

Most people don't consider this option when choosing a career that "makes a difference". But time and money are normally interchangeable – money can pay for people's time, and your time can be used to earn money – so there's no reason to assume the best careers are only those that benefit people directly through the work itself. (MacAskill, 2016, p. 76)

5.4.6 Among the top 10 percent

As already discussed in chapter 4, the richest people in the world live mainly in developed nations. Just to name a few examples: the annual median disposable income per household in the U.S. in 2017 was \$44,110 and CHF 52,000 in Switzerland (OECD, 2020). And those are only the values for the median yearly disposable income. Figure 4 illustrates the global income distribution in 2003 and 2013.




Figure 4 Global income distribution in 2003 and 2013. Graph adopted from Roser (2013)

These numbers have certainly changed somewhat since then, but the trend will certainly have remained the same. According to this, most people living in developed nations earn way above the world mean income and with a median of \$44,110 it is safe to say that a big part of the U.S.-American population belong to the top 10 percent of earners worldwide (Roser, 2013). Individuals earning professional salaries in both nations, are clearly among the 5 percent of global earnings and some even at the top 1 percent.

As MacAskill (2016, p. 78) explained, unlike the typical careers people imagine when thinking about an “ethical” career choice, earning to give can be done by anybody. Certainly, it is a different type of experience and satisfaction than working at one of those highly requested job positions at NPOs or in the public sector. But with a normal job in the developed world, you could maybe have an even bigger impact by allowing those NPOs to hire even more people or expand their programs even further (MacAskill, 2016, p. 78).

5.5 Expected value

What are the chances of success, and how good would success be?

The last main concept of Effective Altruism MacAskill (2016) covered, is the importance of considering the expected value of an action. As the real world is full of uncertainties, in most cases we cannot be absolutely certain the action we choose to perform will really make a difference. Especially as it in most cases of our every-day life, it is almost impossible to know what would have happen otherwise. Among possible scenarios there will always be very few best-case scenarios, which, despite their low probability, if successful, would have a massive impact (MacAskill, 2016, p. 80).

Dealing with uncertainty is certainly one of humanity’s biggest challenges. From trading to political decisions, we face it everywhere. A concept developed to deal with it, is the concept of *expected value*. The idea is to evaluate and compare possible scenarios according to their likelihood and potential adding value. Clearly, in most cases we do not know the exact probability and even the value of each option. But it is still a good framework to work with in the face of uncertainty (MacAskill, 2016, p.81).

5.5.1 Micromort: assessing risk

To illustrate the principle of expected value, MacAskill (2016, p. 82) elaborated on the concept of *micromort*. It is a unit to assess the risk of death of different activities used by public health experts to evaluate and compare the risks of activities ranging from cigarettes to scuba diving. The calculations are based on reported cases of deaths from the activities evaluated. “[...] [O]ne micromort equals a one-in-a-million chance of dying, equivalent to thirty minutes of expected life lost if you’re aged twenty, or fifteen minutes of expected life lost if you’re aged fifty” (MacAskill, 2016, p. 82).

The concept of micromort allows us to make comparisons of relative danger between activities. Scuba diving for example, is five micromorts while flying in a space shuttle is seventeen thousand micromorts, or the equivalent of a 1.7 percent chance of dying. (MacAskill, 2016, p. 82). When evaluating things that might increase our risk of death later in life, such as cigarettes, the effect on life expectancy can also be calculated in minutes. Smoking one cigarette for example, reduces life expectancy by five minutes. Driving an hour in a car reduces our life expectancy by three minutes, while an hour in a train costs you only twenty seconds of life (MacAskill, 2016, p. 83).

The idea of expected value not only allow us to compare the risks of different activities, but it also helps us determine which risks are worth taking. Thinking explicitly in these terms is not very intuitive to most of us because we are bad at assessing low-probability high-value events. Reactions to low-probability effect can be categorized under two extremes, we either give them too much value, e.g., the lottery, or we ignore them altogether (MacAskill, 2016, p. 83).

5.5.2 The Fukushima Daiichi nuclear accident: a black swan

Ignoring low-probability events is what lead to the Fukushima Daiichi nuclear accident in 2011. The nuclear plant was certainly prepared for single natural catastrophes, such as the massive earthquake that first caused the shutdown of the operations of the nuclear reactions. A safety response of the plant to prevent a meltdown. However, they were not prepared to also face the tsunami that followed, which disabled the plant's cooling system and caused the meltdown of three of their reactors. Although only one single person died from radiation-induced cancer, it forced 160,000 people to evacuate their homes, from which 1,600 died during the evacuation (MacAskill, 2016, p. 79).

These are what Nassim Nicholas Taleb calls *black swans* – very low-probability events with massive consequences. They are events like these to which we are just not prepared as no one took the time to think about them before they happened. Only afterwards does the necessity to assess these low-probability risks seem obvious to us (Taleb, 2010).

5.5.3 Expected value: assessing possibilities

However, the concept of expected value cannot only be applied for negative low-probability events, but also for more positive ones, such as voting or eating fewer animal products. There is a very low probability that your vote will be the one deciding the outcome of an election. Depending on where you live, the probability might be higher or lower, but in any case, it is not equal to zero. Your decision to not consume animal products might have similarly low chances to change something. But in the unlikely event, your non-consumption of that packaged chicken might be that one unit remaining at the supermarket that leads the manager of the shop to order less stock in the next purchase, the impact would be huge (MacAskill, 2016, p. 88).

Similarly, the probability of your presence as an individual at a rally making a difference is very low. But, when considering the expected value, if you are the one individual making a difference, it makes sense to go to the rally. This was also the conclusion professors of political science at Harvard and Stockholm Universities made as they analyzed Tea Party rallies held on April's Tax Day 2009. They used the weather as a natural experimental setting. Depending on whether the weather was good, or bad, fewer or more people would show up to a rally. Doing so allowed them to assess whether the number of people at a rally had an effect in the political decision-making. And it did. "Policy was significantly influenced by those rallies that attracted more people, and [...] the larger the rally, the greater the degree to which those protestors' representatives in Congress voted conservatively" (MacAskill, 2016, p. 89).

5.5.4 Expected value: comparing potential careers

According to MacAskill (2016, p. 89) the concept of expected value can also be applied to compare already known ways of doing good with more uncertain strategies that might eventually end up having an even bigger impact. He made a case for expected value allowing us to make comparisons between choosing a highly paid job for doing earning to give and a more difficult but higher-payoff career in politics, research or entrepreneurship (MacAskill, 2016, p. 89). To assess the value of these risky but promising career paths there are several things to consider.

First, the odds of success, which are calculated by looking at the most recent data regarding the career being evaluated. If in politics, what academic background do most politicians working in that sector you are considering have? How much and what type of experience did they gain before starting their political career? How many people fitting these profiles actually seek the same position? (MacAskill, 2016, p. 90f).

Now, with that information in hand, we can continue to estimate the impact they would have in such a career, in order to determine its expected value. To illustrate, MacAskill (2016, pp. 90–94) calculated an example of someone considering becoming a member of parliament (MP) in the UK. Determining the impacts an MP might have is no easy task, as there are many factors to consider, such as restrictions by other interest groups, like other MPs, the electorate or international organizations. These were calculations for the sum of state budget the person might be able to influence, such as allocating spending in overseas development aid instead of defense. He then reduced the value of impact even further, as he considered that even if they would be able to have an influence on the destination of the budget, it would only be a gradual one. This would lead to the budget having a slightly better impact than it otherwise would have had (MacAskill 2016, pp. 90 – 94).

Considering the odds of obtaining the MP position, MacAskill (2016, p. 93) then calculated the expected influence each of the contenders would have. After that, he included the concept of replaceability and reduced the final sum of potentially influenced budget even more.

These are of course all out of the pocket calculations, for which MacAskill applied a method he calls *lower-bound reasoning* (2016, p. 91) – when exact estimations of an action’s impact are impossible to make, he recommends doing estimates we are confident to be an understatement. This way, it is safe to assume that the calculated impact that risky but high payoff action would have *at least* as much impact as that lower-bound estimate.

5.5.5 The expected value of political change

The concept of expected value can, according to MacAskill (2016, p. 94) also be used to try to assess the value of trying to influence political change. Clearly, the exact probability and even the impact of systemic change are very hard to assess, but the potential benefits are so big that rough estimates usually suffice. He even argued that supporting the right systemic change causes might have even greater effects than donating to the most effective organizations (MacAskill, 2016, p. 94). So, “if you can find the right area, funding or participating in political campaigns could potentially do even more good. [...] The point is simply that long shots can be worth it if the payoff is big enough” (MacAskill, 2016, p. 94).

5.5.6 Why should the status quo be the end of political change?

When learning about potential courses of action to change the political system, MacAskill (2016, p. 94) made the case that effective altruists should not discard such ideas as inefficient just because they seem unlikely to happen. In fact, many ethical ideas we regard today as commonsense were perceived as highly radical in the past. We just have to think about the historically very recently obtained equal rights for women, black people, indigenous people, and non-hetero sexual people to understand this.

Benjamin Franklin (who himself owned slaves) wrote an open letter to the first US Congress petitioning the abolition of slavery, causing resignation and outrage among the defenders of slavery (National Archives, 2016). Questions such as ““Who is going to compensate the slave owners?” [...] and, “What is the mixing of races going to do to American values and character?”” were raised during the debate, unthinkable for today’s social standards (MacAskill, 2016, p. 94).

Jeremy Bentham, one the first figures of utilitarianism, spoke early on about the ethical dilemmas of several systemic issues in society, many of which we have not solved yet completely. He was for the abolishment of slavery (Crimmins, 2019) and the death penalty (Bedau, 1983), and demanded the liberation and suffrage of women (Williford, 1975). Bentham generally addressed several then taboo and controversial aspects of society at his time. He was in favor of sexual liberation and criticized the punishments in place in the UK for then “unwanted” sexual behaviors, such as sodomy, masturbation and even homosexuality (Crompton, 1978). He was also the first philosopher to speak out against the suffering of non-human animals (Kniess, 2019).

Certainly, these are all positions most citizens in the Western world share today. But they were very counterintuitive for the big majority of our predecessors. The abolition of slavery, women suffrage, LGBTQ+ rights are historically speaking very recent events. But these positions and the legitimacy of these groups were not only questioned but systematically oppressed in the past. Their chances for success were most probably assessed as very low, but the benefits they have brought to those communities are undeniable today. So, why should the status quo be the end of political change?

5.5.7 Low-probability high-impact risks

The concept of expected value can also be applied when it comes to the current climate crisis. The discussion on whether climate change is actually man-made or not becomes irrelevant when viewed through the concept of expected value. “*Even if scientists had not already shown that man-made climate change is happening, the mere fact that man-made climate change might be happening is enough to warrant action*” (MacAskill, 2016, p. 95). The potential consequences of climate change are just too big to take the risks. Just to make a recent and even more palpable analogy: the risk of a new virus developing into a pandemic is very small, but as we can experience ourselves today, the consequences are just immense. These are topics further discussed below in chapter 6.2.

To sum up the essence of expected value, MacAskill summarized the importance of handling these types of low-probability high-impact scenarios as follows:

Just as most of the value from aid programs comes from the very best aid programs [...], it's often the case that most of the expected harm from disasters come from the very worst disasters. (That is: the death tolls from disasters form a fat-tailed distribution. [...]). (MacAskill, 2016, p. 98)

So, even if most of the examples and concepts discussed above were made around hard data, like deworming schoolchildren, Effective Altruism also addresses issues that are harder to grasp. Seemingly “unquantifiable” areas like political change or low-probability high-impact risks are also issues that effective altruists consider relevant to discuss when it comes to “making the world a better place”. Assessing the expected value in these areas is certainly hard to do. But MacAskill (2016, p. 99) argued that we should at least have broad estimations for us to be able to make better decisions.

5.6 In the classroom

Up until this point the class should already be familiar with Effective Altruism, its definition(s), the thought experiment and the ethical and logical arguments behind the concept. The pupils should also be aware that they will very probably belong to the richest 10 percent globally. The next step, discussed in this chapter, is to think about and ideally discuss how they would like to set their limited resources, striving to achieve the greatest good possible. The Socratic dialogue could certainly be useful during these discussions. For this, the teacher needs to set the right questions, that should ideally guide the students through a (joint) reflection towards a consistent answer.

This process has many advantages, at the motivational, it encourages pupils to involve (through higher intrinsic motivation) with the question and gives the class a touch of gamification at the cognitive level presenting the topic as a solvable enigma. The Socratic dialogue could be especially useful when discussing these aspects of EA, as it enables the teacher to guide the class through the five key elements of EA – scale, effectivity, neglectedness, alternative and expected value – without the students necessarily knowing all the facts discussed above; logical thinking would suffice to understand that, e.g., one intervention, *ceteris paribus*, is superior to another, if it helps more individuals or if it helps the same amount of individuals better (scale).

Students would not only learn to think how resources should be set to make the world a better place, but the key aspects discussed above are useful (scientific) thinking tools for better decision making in all areas of life. This is certainly another good argument, beyond the ethical aspect, to discuss Effective Altruism with pupils in classrooms. How exactly I propose to present these principles in class is discussed in one of the lesson plans presented in Chapter 8.

6 Effective Altruism's biggest cause areas

Having discussed what Effective Altruism is and some of its main concepts, this chapter gives an overview on the current subjects – known as cause areas – most effective altruists are currently supporting in one way or another. The focus of this chapter lies in the latest results of the Effective Altruism Survey from 2019. The last section covers further cause areas, that do not take the center stage, but are still addressed by some members of the community. Now, what exactly is a *cause area*?

A “cause” is the field around a particular problem or opportunity — such as reforming the criminal justice system, preventing pandemics, or reducing the burden of Alzheimer’s disease — in which it’s necessary to develop expertise and networks to make good giving decisions. (Open Philanthropy, 2015)

This chapter will further expand on the top four cause areas, to help the readers understand the reasons for effective altruists to consider these causes as the most relevant problems of all. Clearly, this is only a snapshot of the status quo; the focus of the movement might change as members of the community change their opinions and or new cause areas arise. As described above, the Effective Altruism community is always considering the possibility of having overlooked problems (EffectiveAltruism.org, 2020) and they do not pretend to have evaluated all potential problem areas relevant for EA (Todd, 2014). There is even a whole cause area dedicated to finding new cause areas, called Cause X (Joey, 2019). Discussions on cause prioritization among community members can be followed among others in the Effective Altruism Forum (EA Forum, n.d.).

6.1 Shifting focus in the Effective Altruism community

When reading through the EA Forum, a shift in the cause prioritization of the EA community away from poverty alleviation towards the long-term future and the reduction of catastrophic and existential risks becomes clear. This was also in the results of the latest survey of the Effective Altruism community in 2019, where 41 percent of participants chose catastrophic and existential risks reduction when pressured to choose only one cause. Many of the respondents (42 percent) said they have changed their main cause area – the cause area they mainly focus on – since joining the movement, and most of them (54 percent) have moved from other cause areas towards catastrophic and existential risk reduction (Dullaghan, 2019).

Indeed, Peter Singer warned in 2019 that the shift towards existential risk as the public face of EA might be problematic, due to its inaccessibility to the broader public (Wiblin et al., 2019). Despite this, global poverty remains the most popular single cause among effective altruists, as illustrated in Figure 5 with 22 percent of participants.




Figure 5 Comparison between 2018 and 2019 responses in the Effective Altruism community survey when participants were asked to choose one cause to be the top priority in Effective Altruism. Graph adopted from Dullaghan (2020).

The next two most popular causes among participants of the survey were climate change (16 percent) and artificial intelligence (AI) risk (14 percent), which are fields of the cause area focused on reducing catastrophic and existential risks, also known as long term future. Eleven

percent chose cause prioritization as their top priority. Animal Welfare was highlighted as the top priority cause by eight percent of respondents, exactly as many in favor of improving decision making to be top priority. The rest of respondents distributed their answers between Other X-risks, Meta charities, with five percent each, Biosecurity with four percent, and Mental Health and Nuclear Security with three percent each (Dullaghan, 2019).

In comparison to the survey results of 2018, there was a significant increase in respondents favoring climate change as the cause area to be prioritized (Dullaghan, 2019). Clearly, the survey only provides a peak view into the minds of a fraction of the Effective Altruism community. Without longitudinal panel data it is impossible to determine whether this pattern is due to under- or overrepresentation, members of the community changing their opinions or due to new members who happen to prioritize climate change in particular (Dullaghan, 2019).

However, it would be wrong to assume that respondents only cared about the cause area they selected as their top priority. When asked about their opinions on each of the listed cause areas, a very different picture arises. “Most listed causes received substantial support, with no cause receiving fewer than 59% of respondents suggesting that it should receive at least “significant resources”” (Dullaghan, 2019). The top five cause areas, as ranked by the respondents of the survey, were Global Poverty, Cause Prioritization, AI Risks, Climate Change, and Biosecurity. All of them were indicated by at least 44 percent of respondents to be near-top or top priority. Indeed, all but two causes (Mental health and Nuclear security) were judged by ~34-49 percent of participants to be near the top priority or top priority, while 62 percent of respondents judged Global poverty as top or near top priority (Dullaghan, 2019). These figures on cause prioritization coincide with the latest survey of some Effective Altruism organizations leaders (Gertler, 2019) and 80,000 Hours’ problem profiles page (80,000 Hours, n.d.).




Figure 6 Likert graph depicting cause areas listed in descending order based on the average of their full scale ratings from the Effective Altruism survey 2019. Graph adopted from Dullaghan (2020).

In the next sub-chapters provide a short overview of the top four cause areas – Global poverty, Reduction of catastrophic and existential risks, Cause prioritization and Animal welfare. This includes also a short listing of organizations currently working to address the problems.

6.2 Global poverty

As described above, the philosophy of Effective Altruism started by highlighting the ethical dilemma we face when it comes to the struggles of people living in (extreme) poverty around the world and the various possibilities we have to help reduce it. So, it is no wonder, that global poverty remains the top priority among effective altruists. This chapter gives a short introduction into the scale and urgency of poverty and provides the readers with a small overview of organizations currently helping to reduce it.

A widely accepted definition of extreme poverty is that by the World Bank, which defines it in absolute terms. Extreme poverty is defined as living on less than \$1.90 per day, and moderate poverty as living on less than \$3.10 per day. Both figures were adjusted by the World Bank in purchasing power, meaning they represent the “[...] daily consumption of goods and services – whether earned or home-grown – comparable to the amount of goods and services that can be bought in the United States for \$1.25 [the extreme poverty line in 2010, today at \$1.90].” (Singer, 2010, p. 8 comment from The World Bank, n.d.-c).

In 2019, there were 690 million people living under extreme poverty. This number, as already described above, is expected to rise significantly due to the Covid-19 pandemic. Most of the individuals afflicted by poverty and extreme poverty are women and children (The World Bank, n.d.-c).

6.2.1 Extreme poverty is the breeding ground for extreme suffering

To understand what it means to live under (extreme) poverty, this section provides the reader with a small comparison between affluent and poor nations of some common denominators for comparing well-being, known as world development indicators. The stark contrast in living standards between developing and developed countries is still immense. Life expectancy for people in the European Union and Canada is up to 25 years longer than for someone in Nigeria or Sierra Leone. The difference remains high even compared to somewhat more developed nations like India or Bolivia, where people's life expectancy is ten to twelve years shorter (Roser et al., 2019).

Child mortality in Chad or Central Africa is at least 27 times higher than in the European Union (Roser et al., 2019). For those aged between 15 and 49 in developing countries, among the most common causes of death are infectious diseases, like HIV/AIDS and tuberculosis. Even diarrheal diseases and malaria are among the most common causes of death for adults in the poorest nations worldwide. While people in richer countries such as in the case of Germany, die mostly from cancer, cardiovascular diseases and suicide (Ritchie & Roser, 2019).

Maternal mortality has decreased all around the world but in Sub-Saharan Africa, where two thirds of all maternal deaths occur. South Asia has the second highest rate, at 22 percent. Clearly, this is in part due to the high fertility rates in Asia and Africa, but they also lack key medical care during pregnancy and birth which save women's life in rich nations. In Nigeria, for example, a woman is 200 times more likely to die in pregnancy or childbirth per birth than a woman in Sweden (Roser & Ritchie, 2013).


Public health professionals use, besides the QALY, the concept of the disability-adjusted life years (DALYs). This value, measuring the overall disease burden includes both, mortality and morbidity, quantifies the years of "healthy" life lost due to disease. Its values are usually given per 100,000 individuals. While Canada, France and South Korea have values below 20,000

DALYs, the values of many countries in Sub-Saharan Africa are higher than 80,000 DALYs (Roser & Ritchie, 2016) as can be seen in Figure 7.

Burden of disease, 2017

Disability-Adjusted Life Years (DALYs) per 100,000 individuals from all causes.

DALYs measure the total burden of disease – both from years of life lost due to premature death and years lived with a disability. One DALY equals one lost year of healthy life.



Source: IHME, Global Burden of Disease

Note: To allow comparisons between countries and over time this metric is age-standardized.

OurWorldInData.org/burden-of-disease • CC BY

Figure 7 World map of Disability-Adjusted Life Years (DALYs) per 100,000 individuals from all causes. DALYs measure the total burden of disease - both from years of life lost due to premature death and years lived with disability. One DALY equals one lost year of healthy live. Figure adopted from Roser & Ritchie (2016)

6.2.2 Covid19: the first progress setback in 20 years

In the last decades, we had actually made immense progress in the fight against extreme poverty, reducing it from 42.5 per cent of the world population in 1981 to 9.2 per cent in 2017 (The World Bank, n.d.-b). This number, however, is expected to grow due to the current Covid-19 pandemic, causing the first setback in poverty reduction in 20 years (The World Bank, n.d.-a). According to preliminary estimates, “[...] an additional 88 million to 115 million people will be pushed into extreme poverty, bringing the total to between 703 and 729 million” (The World Bank, n.d.-c).

6.2.3 Poverty hits the most vulnerable segments of the population harder

Especially women and children will be mostly affected by this reversal. Even in 2019 around 5.2 million children under five years died from mostly preventable and treatable causes such as birth complications, pneumonia, diarrhea and malaria. All of these diseases are easily preventable or adequately treatable with simple and inexpensive interventions such as adequate nutrition and hydration as well as immunization (World Health Organization, 2020) and the distribution of bed nets. Figure 8 illustrates the number of deaths by age at a global level in 2017, a year in which more children under five died than people between 75 – 79 years (Roser et al., 2019).

Deaths globally by age

56 Million people died in 2017. Shown here is at what age each person died.

Our World
in Data




Figure 8 Deaths globally by age in 2017. Graph adopted from Roser et al. (2019)

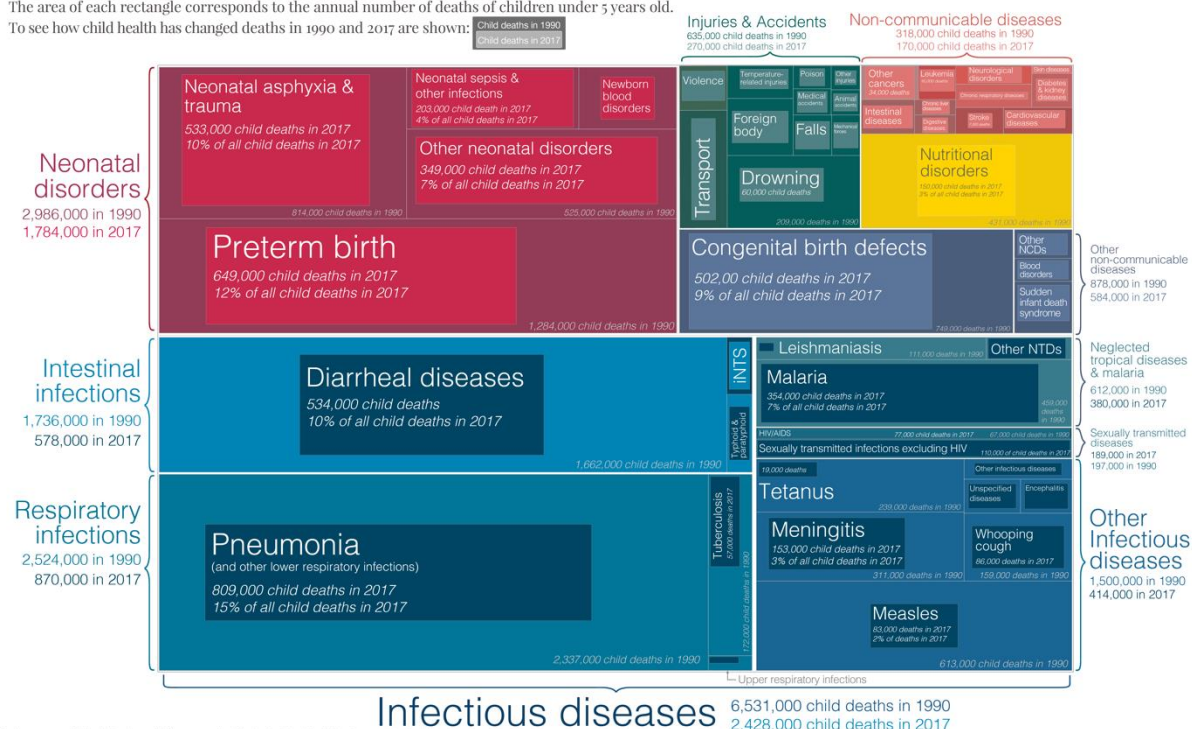
The reasons for their extremely premature deaths are diseases and disorders, people in the West have long overcome. Figure 9 gives an overview of the main causes of death of children under 5 years of age. Despite child deaths having more than halved since the 1990s, from 11.8 million to 5.2 million in 2019, there are still too many children whose death we could have easily prevented. Pneumonia and other lower respiratory diseases due to bacterial infections are the cause of death for one in seven children. Prematurely born children in poor countries do not receive the attention and care they need, prematurely born children in affluent nations certainly

get. This lack of care is also the reason why congenital defects, such as physical or genetic abnormalities end up being lethal for the children affected. But far worse, is the fact that diarrheal diseases, an easily treatable condition, are the third leading cause of child mortality (Roser et al., 2019).

What do children die from? And how have the causes of child death changed since 1990?

Our World
in Data

The area of each rectangle corresponds to the annual number of deaths of children under 5 years old.
To see how child health has changed since 1990 and 2017 are shown:



Data source: Global Burden of Disease - Institute for Health Metrics and Evaluation.
OurWorldinData.org - Research and data to make progress against the world's largest problems. Licensed under CC-BY by the author Bernadeta Dadonaite.

Figure 9 Chart of the causes of death for children under 5 years old. The area of each rectangle corresponds to the annual number of deaths of children under 5 years old. Graph adopted from Dadonaite (2019)

Almost half of all child deaths are caused by infectious diseases. Many such as tuberculosis and hepatitis can be prevented through immunization or in the case of bacterial infections, treated with antibiotics. In the case of malaria, we know that the best way to protect children is preventing them from getting it in the first place through insecticide treated bednets. That, at least until the new malaria vaccine programme has been set up (Roser et al., 2019; WHO, 2019a). Malaria also counts as one of the most common causes of death for children aged between 5 and 14 in poor countries. Other common reasons in this age group in developing countries are HIV/AIDS and to a lower extent diarrheal diseases and lower respiratory infections (Ritchie & Roser, 2019).

Clearly, there are many more factors making life more difficult for poor people. Those that survive childhood, are often malnourished and have long-lasting negative effects if they felt sick during childhood. Children who grew up malaria free earn 50 percent more per year, during their entire life (Banerjee & Duflo, 2012, p. 45). We now understand that saving money is very difficult for poor people as they live from day to day. And without savings they can easily get into debt, especially if a family member falls sick, as the only loans available to them have extremely high interest rates (Banerjee & Duflo, 2012, pp. 43–45).

But not only are there physical health factors, but poorness often comes hand in hand with a degrading state of powerlessness. Many poor people experience a range of situations in which they have had to endure humiliation or even neglectedness by their state's police (Singer, 2010, p. 6).

6.2.4 Organizations working on this cause area

In the chapters above, when discussing the key concepts of Effective Altruism, I already named a few organizations related to Effective Altruism and further recommended by the community. Maybe the most salient organization is GiveWell, a meta-charity, dedicated to evaluating the work of other organizations in their effectiveness. “[They] search for the charities that save or improve lives the most per dollar” (GiveWell, n.d.-a). For this they assess organizations along four criteria: Their evidence of effectiveness, cost-effectiveness, transparency, and room for more funding (GiveWell, 2016). Every year, they publish an updated review of the organizations that would make the most out of new donations, according to their criteria.

In 2019 alone, GiveWell was able to influence at least \$152 million – the sum they were able to track, the grey numbers might certainly be higher than that. This is already more than \$10 million more than in 2018. The sum entails the money moved to the charities they recommend, including the grants the Open Philanthropy – a NGO dedicated to research potential focus areas and investigating giving opportunities – handed out to GiveWell's recommended charities. GiveWell's reported sum also entails the money moved via the Incubation Grants program (GiveWell, 2020e).

For their evaluations, they regularly contact the organizations they review for updates on their latest inputs and funding, to re-check their evaluations and make sure the program has room for more funding. Their reviews of all the organizations they have contacted until now as well as those they are currently in contact with or plan to do so are accessible on their website (GiveWell, n.d.-b). All of their evaluations are open access, accessible to anyone who wants to read them, not only for transparency but they also actively welcome criticism on their approach and calculations. They also collect the mistakes they have made in the past and what they have learned from them on their website (GiveWell, 2020a).

GiveWell's strongest recommended organizations, also called top charities, work, as already described above, primarily in poor countries. The top two charities are both fighting malaria in the developing world, a disease that kills over 400,000 people every year, and mostly children under five years of age in sub-Saharan Africa (GiveWell, 2020d). There are two charities GiveWell recommends in this area: Malaria Consortium – Seasonal Malaria Chemoprevention and Against Malaria Foundation. The first distributes preventive antimalarial drugs for children between three and 59 months old in four monthly cycles during high transmission season. This helps prevent illness and death from malaria in babies and small children, those most affected by the disease, at \$7 per child (GiveWell, 2020c, 2020d). The Against Malaria Foundation is dedicated to the distribution of long-lasting insecticide-treated nets in developing countries at about \$5 each. GiveWell was able to track both organizations receiving almost \$53,000,000 and \$30,000,000 respectively, through GiveWell's outreach in 2019 (GiveWell, 2020e).

The next top charity is Helen Keller International, more specifically it is a vitamin A supplementation program. As advocates and providers of technical assistance and funding to government-led vitamin A supplementation programs in developing countries they are trying to reduce deaths due to Vitamin A deficiency, which kills more than 200,000 children every year. The costs are minimal, at \$1 per vitamin A supplement (GiveWell, 2020d). The next top organization, New Incentives is dedicated to promoting the immunization of infants in Nigeria. As parents lose at least part of their daily income if they take their child to the clinic for immunization, they are prone to postpone it indefinitely. This has led to 43 percent of infants in Nigeria not receiving all recommended childhood vaccines. GiveWell argues that through the program of New Incentives, it takes about \$47 to incentivize caregivers to complete a baby's immunization program (GiveWell, 2020d).

As already mentioned in the chapters above, Evidence Action's Deworm the World Initiative is also regarded as one of the most cost-effective interventions to fund. "Donating \$100 will deworm about 100 children, leading to a cumulative \$1,100 in additional earnings over the course of the children's lives" (GiveWell, 2020d). The program is run in regions where the average person lives on \$700 per year. Besides the Evidence Action's initiative, there are other organizations dedicated to do the same that have been assessed to be similarly effective. These are SCI Foundation, the END Fund's Deworming Program and Sightsavers' Deworming Program (GiveWell, 2020d).

Another recommended charity is GiveDirectly, which gives \$83 out of every \$100 donated in cash to poor families in poor nations, primarily in Africa, to spend as they like. Many of the beneficiaries make investments in housing, education and business and agricultural assets. GiveWell also gives donors the option to donate to their Maximum Impact Fund, whose final destination they decide based on the latest evidence they obtain on a quarterly basis. They do not take any fees for this service (GiveWell, 2020d).

Besides GiveWell, The Life You Can Save (TLYCS), an organization named after Peter Singer's book, encourages people to pledge a fraction of their income to effective charities. They also recommend GiveWell's recommended charities among others (TLYCS, 2021). Good Ventures, a philanthropic foundation also based in the USA, cooperates with GiveWell and Open Philanthropy in promoting GiveWell's top charities and grants funds to (new) promising initiatives in line with Effective Altruism values (Good Ventures, n.d.).

Similar to GiveWell, Giving What We Can, already named in chapter 1, evaluates charities as well and encourages people to pledge 10 percent of their income to effective charities. However, their charity recommendations are not limited to organizations helping the global poor, but cover other areas as well, such as non-human animal welfare and long-term future. They also recommend GiveWell's Maximum Impact Fund among other funds for the other biggest cause areas of Effective Altruism for donations (GWWC, 2021).

The last organization related to Effective Altruism and global poverty worth noting, is AidGrade. They evaluate the cost effectiveness of poverty reduction measures but concentrate thereby less on individual organizations and more on evaluating the intervention per se. For this, AidGrade focuses primarily on randomized controlled trials to assess the effects of a given program (AidGrade, n.d.).

6.3 Reduction of catastrophic and existential risks

As the Covid-19 pandemic shows, there are low-probability events that might result having dire consequences for the whole of humanity. The emergence of a corona virus strain able to infect humans at one life market in Wuhan, has led to one of the worst public health crises and certainly the worst financial crisis we have experienced since WWII. It is risks like this, that the scientific community and also the Effective Altruism movement has been warning us about for years. Effective altruists argue that artificial intelligence, self-replicating nanotechnology and superbugs are all low-probability risks, that exactly as the risk of a pandemic, have not received the attention they deserve.

Besides caring about our own well-being, many people care about the wellbeing of future generations. The number of people who could exist in the future is probably many times greater than the current number of living people. Just as we should care about the people, living on the other side of the globe, the Effective Altruism community argues that the lives of the people living in the future should count as well. Due to this, keeping the world a safe place for human beings to live has also become one of the key areas of focus for people around Effective Altruism (EffectiveAltruism.org, 2020).

6.3.1 Risks with varying probability

There are many ways the quality of our and future lives might get worse. Among the most well-known risks are human-induced climate change, nuclear war and due to current circumstances, pandemics. However, despite the high levels of public awareness and the public demand to act, especially in regard to the climate crisis, we are very far from doing enough to avoid the worst from happening. The number of resources currently invested in their mitigation (and eventual management) is, as we now know with respect to pandemics, too small in relation to the destruction they (might) cause.

These are by no means the only risks experts have been warning us about. Many scientists worry about the risk future technologies, such as advanced artificial intelligence, bioweapons but also self-replicating nanobots and synthetic biology might impose on us. One of the most prominent figures speaking about these risks is Nick Bostrom, University Professor at Oxford and author of *Global Catastrophic Risks* (2008) and *Superintelligence* (2014), who has dedicated his life to the understanding and mitigation of such risks.

Thoughts about existential risks – events that might end up causing the extinction of the human species – can be traced back to the 19th century. But the academic field of Existential Risk Studies (ERS) as such only took off at the beginnings of the 2000s – in part thanks to Effective Altruism since 2009 (Beard & Torres, 2020, p. 1f). Today, “[t]here is an emerging scientific consensus that, due to the multiplicity of risks with the potential to cause global catastrophes, homo sapiens is now in the most perilous moment of its 300,000- year history” (Beard & Torres, 2020, p. 2). Authors argued that the challenges we are currently facing are of such magnitudes, that all catastrophes and crisis we have experienced to this day, are mere “ripples on the surface of the great sea of life” (Bostrom, 2002, as cited by Beard & Torres, 2020, p.2).

With this in mind, Bostrom and others argue that the prevention of existential risks should be a global priority. “Many theories of value imply that even relatively small reductions in net existential risk have enormous expected value” (Bostrom, 2013, p. 15). As already depicted above, effective altruists consider this cause area to be at least among the top cause areas of Effective Altruism, especially so in regard to the dangers of climate change and artificial intelligence (AI).

Certainly, efforts towards this end have improved severely since 2009; organizations such as the Future of Humanity Institute (FHI), the Global Catastrophic Risk Institute (GCRI), the Future of Life Institute (FLI) and the Centre for the Study of Existential Risk (CSER) have been established to further ERS. However, most experts still think that we are far from having a good understanding of the risks and lack the necessary precautions and mitigation measures to deal with them in case they occur (Todd, 2017). In the next sections, I briefly addressed two of the biggest risks discussed in ERS: pandemics, superbugs and artificial intelligence.

6.3.2 The risk of a pandemic

Researchers all around the world as well as the World Health Organization (WHO) have been warning us for decades about the potential development of new strains of viruses and bacteria resistant to our current medical technology (Henig, 2020). In fact, just months before the Covid-19 pandemic broke out, the WHO published its last report on the risks of a global pandemic. In it they criticized governments around the world for failing to prepare for such a scenario (Global Preparedness Monitoring Board, 2019); all while the first Covid-19 strain was just about to emerge in Wuhan. In their foreword, the WHO summarized the case for a global system of preparedness as follows:

The world needs to proactively establish the systems and engagement needed to detect and control potential disease outbreaks. These acts of preparedness are a global public good that must meaningfully engage communities, from the local to the international, in preparedness, detection, response and recovery. Investing in health emergency preparedness will improve health outcomes, build community trust and reduce poverty, thereby also contributing to efforts to achieve the United Nations Sustainable Development Goals. (Global Preparedness Monitoring Board, 2019, p. 6)

The current Covid-19 pandemic is the most vivid examples of the dangers they have been warning us about. In just a matter of months, the corona virus had reached every single country in the world and filled up intensive care units at hospitals all around the world in a matter of weeks (Our World in Data, 2021a). It brought the world economy to a stop, led to unemployment levels not seen in decades, and has caused the death of almost 3 million people worldwide and counting (Our World in Data, 2021b).

Our fight against both viral and bacterial germs has resulted in an evolutionary treadmill, as we create the perfect circumstances for a high-speed version of natural selection. This is already resulting in the development of new viral strains for which our current vaccines are no use for and new microbial strains resistant to multiple antibiotics, also known as *superbugs*. This process is not only taking place at hospitals around the world (Cornejo-Juárez et al., 2015; Matta et al., 2018) but primarily in the global food industry through intensive animal farming. These systems, broadly known as factory farming have developed as the most efficient way to raise billions of animals to satisfy our lust for meat and other animal products. Today, 90 percent of animals worldwide are held this way. The crumbed circumstances the animals live in create the perfect environment for the development and quick spread of diseases. To avoid

this, every single animal would need to be vaccinated properly and cases of bacterial infections would need to be handled adequately without overusing antibiotics. However, the reality looks more like this:

[A]nimals on factory farms are routinely fed vast amounts of antibiotics in order to keep them alive in conditions that would otherwise kill them. Because of this, even the most powerful antibiotics aren't effective against certain bacteria, contributing to the emergence of "superbugs" – new, aggressive, antibiotic-resistant pathogens. (Chittilapally, 2020)

It is common practice in the animal farming industry to constantly make use of antimicrobials in low doses with the goal of preventing the diseases from developing in the first place (Weissman, 2015; You & Silbergeld, 2014). Currently in the United States, 80 per cent of the nation's annual antibiotic consumption is implemented in livestock (Food and Drug Administration, 2014). This includes many antimicrobials that are also important in human medicine for the treatment of common infections as well as during major surgeries, organ transplantations, and chemotherapy (Laxminarayan et al., 2013). This malpractice of antibiotics in livestock has been linked multiple times to the development of multi-resistant bacteria with massive public health risks (You & Silbergeld, 2014). In the US alone, multi-resistant bacteria cause the premature death of 35,000 people each year (Chittilapally, 2020). However, it is not only the artificial selection of bacteria resistant to multiple antibiotics that should worry us, but also the accelerated development of new viral strains.

The Severe Acute Respiratory Syndrome (SARS) outbreak in 2002-2004 caused by the first identified strain of SARS coronavirus species (SARS-CoV-1) had a fatality rate of 11 per cent. Similarly, the swine flu (H1N1) pandemic in 2009 affected at least 11 per cent of the global population and killed between 151,700 and 575,400 people, primarily in Southeast Asia and Africa. It comes as no surprise that before 2009 there had already been positive cases of human infections from new swine flu strains, however all mainly asymptomatic until H1N1 emerged in 2009 (CDC, 2019c).

The Middle East Respiratory Syndrome (MERS) caused by a coronavirus as well had a fatality rate of 30-40 percent. First reported in Saudi Arabia in September 2012 the illness seems to have first occurred in Jordan in April 2012 and stem from dromedary camels (CDC, 2019a, 2019b). The largest registered outbreak of MERS outside the Arabian Peninsula was in the Republic of Korea in 2015, due to one traveler returned from the Arabian Peninsula (CDC, 2019a). MERS-CoV affected more than 2,400 people worldwide and had a fatality rate of 35 percent (WHO, 2019b).

New epizootic virus strains such as variations of the avian influenza affecting chickens, emerge regularly and end up spreading around the world in factory farms. The H5N1 strain from 2005 is one such example. Despite them not being dangerous to humans (yet), they cause the death of millions of chickens who are killed in mass graves to reduce the spread of the disease. And the virus might very well affect us in the future as well. As zoonotic diseases, which make about 60 per cent of all known infectious diseases, the development of the next virus transmittable to humans as well is only a matter of time (Chittilapally, 2020). Furthermore, contamination with multi-resistant bacteria is extremely easy, as humans can catch them either via the environment, food products or direct contact with the sick animals (Boeckel et al., 2015). On February 2021, the Russian Federation notified the WHO of having detected avian influenza A in seven poultry farm workers – the first known animal to human transmission of avian influenza A (WHO, 2021).

6.3.3 Artificial Intelligence

[As] Vladimir Putin said: “*Artificial intelligence is the future, not only for Russia, but for all humankind. It comes with enormous opportunities, but also threats that are difficult to predict. Whoever becomes the leader in this sphere will become the ruler of the world.*” (Vincent, 2017, as cited by Moses & Sheryl, 2019, p. 9)

By the beginnings of the century, experts had already identified several threats to humanity, artificial intelligence being one of the most prominent ones (Good, 1966). The international race towards *artificial general intelligence* (AGI) – “general-purpose systems with intelligence comparable to that of the human mind (and perhaps ultimately well beyond human general intelligence)” (AGI Society, n.d.) – worries scientists around the world about the potential negative consequences the development of AI might bring with it. Indeed many of them think that there is a greater than 50 percent chance of high-level machine intelligence being created by 2050 (Open Philanthropy, 2016). By discussing and collaborating in the development of super-intelligence, they strive for the development of a safe and globally beneficial AI in both the near- and long-term.

Most current AI systems can be understood as *narrow* applications of AI – systems that are superior to human performance in a well-specified field, such as a particular game. Deep Blue’s victory against chess world champion Gary Kasparov in 1996/7 was a historic moment, as the first defeat of human intelligence by a machine. Today’s translations, search engines, image and text generation and recognition as well as social media algorithms are all examples for narrow AI (Piper, 2018).

The first less narrow AI was DeepMind; given only the order to maximize the score in Atari, it was able to reach scores no human had ever achieved beat in all Atari games (Walker, 2020). In 2016, AlphaZero, a generalized version of AlphaGo, mastered chess, Shogi and Go in only a matter of hours via simulations. And most recently, AI mastered an e-sport game for the first time in history as AlphaStar made it to the top 0.02 percent of players in Starcraft, a multiplayer strategic game, it taught itself to play (Walker, 2020).

DeepMind, AlphaZero and AlphaStar are examples of today’s growing less narrow AIs, all of which (literally) learned the tasks they were given via simulations – an AI technique known as *deep learning*. It is these types of generalized learning capabilities, in which computers teach themselves how to achieve a goal, that has put the development of artificial general intelligence (AGI) in the fast tracks (Piper, 2018).

At the same time, we are now realizing that narrow AIs eventually develop abilities that no longer fit the definition of narrow AI, such as GPT-3. It is the latest text-generating system from OpenAI, which has been described as uncanny and fascinating by those who have used it (Piper, 2020).

This massive progress in AI raises several ethical questions. First, the data AI systems are trained on is not necessarily impartial. Social media algorithms are designed to keep us on their sites as long as possible. This seems to not only have an effect on our mental health (O'Reilly et al., 2018) but is in part to blame for our societal polarization and the rise of conspiracy theories in society (Banks et al., 2020). The impact Cambridge Analytica and associates had in the 2016 U.S. election and in U.K.'s Brexit referendum are just examples of the level of social manipulation "simple" AI is capable of (Isaak & Hanna, 2018). AI is also being applied by the government to "optimize" societal systems, e.g., determine how likely convicted felons are to reoffend. Here, the data used to train the algorithms is biased against black and poor people (Piper, 2018).

These problems are small examples of the main problem researchers have been warning us about: the *AI alignment problem* – "[...] how to ensure that the instructions we give to a machine intelligence really capture what we want it to do, without any such unintended outcomes." (Wiblin, 2017a). The smarter a system gets and the more possibilities of action it has, the more critical the input we give it becomes. It is important to understand that intelligence does not include values and morals, so we need to be extra careful not only with the data we feed it but also with the goals and conditions we set. Many experts and effective altruists think that solving this *control problem* – making sure AI does what we want it to do without unexpected "side effects" – should be our priority number one (Wiblin, 2017a). As Wiblin (2017a) said:

It's hard to imagine a more important research question. Solving the control problem could mean the difference between enormous wealth, happiness and health — and the destruction of the very conditions which allow humanity to thrive. (Wiblin, 2017a)

In recent years there has been an increasing number of organizations dedicated to this goal. The Machine Intelligence Research Institute (MIRI), for example, campaigns for *transparent* AI approaches, that is, precisely specified decision algorithms rather than "genetic" algorithms.

This would avoid the development of algorithms whose behavior humans are not able to comprehend. Their research focuses thus in “[...] putting theoretical foundations under AI robustness work” (MIRI, n.d.).

The Future of Humanity Institute (FHI) is a research institute at the University of Oxford, led by Founding Director Professor Nick Bostrom. It combines mathematics, philosophy and social sciences to address pressing questions about humanity. At its core is the Centre for the Governance of AI, which deals exactly with the question of AI alignment and AI safety (FHI, n.d.). Similarly, the Global Catastrophic Risk Institute (GCRI) is a nonprofit think tank with the goal of assessing and developing management strategies for potential catastrophic and existential risks (GCRI, n.d.). The Future of Life Institute (FLI) and the Centre for the Study of Existential Risk (CSER) are further examples of organizations dedicated to the subject of catastrophic and existential risks. Google DeepMind and OpenAI are for-profit organizations that are also working on the subject.

Total spending in AI safety reached \$9 million in 2017, 100 times less than the total spending in the development of machine intelligence in the private sector. Scientists argue that this is still far from enough to properly address the problem. Due to this level of neglectedness, effective altruists think that investing more resources into this problem would lead to rapid progress in the field (Wiblin, 2017a).

6.4 Cause prioritization

Resources are limited but problems are vast. As already discussed above when describing Effective Altruism, it is important to understand that when we make a decision to put our resources in one area, we are automatically deciding to invest less in other areas. Governments, foundations and individuals together spend over \$500 billion on efforts to improve the world. Now, the question is, which problems should these resources be directed to: immigration policy, research, international development, or something else?

Despite the importance of deciding where to invest this massive sum, less than 0.01 percent of it is spent in global priorities research, ergo, determining which global problems are the most pressing to work on. Influencing where these funds end up is considered an immense opportunity to do good in Effective Altruism. Practically, this is the macro-scale application of Effective Altruism to the global problems; determining which are the most pressing issues, looking for cause areas that might have been overlooked until now or remain neglected despite their urgency, and evaluating the results of the interventions funded until now (Duda, 2018).

In fact, Effective Altruism has already been successful in redirecting hundreds of millions of dollars. GiveWell alone was able to influence the destination of more than \$152 millions in 2019 alone (GiveWell, 2020e). Open Philanthropy recommended over \$200 million worth of grants in the same year in projects ranging from X risks and farm animal welfare to scientific research and criminal justice reform (Open Philanthropy, 2021). The Global Priorities Project (GPP) was able to help reallocate £2.5bn of the UK Department for International Development to fund research into the treatment and response to the diseases that cause the most suffering, such as malaria. The U.S. Bush administration launched the \$1.2bn President's Malaria Initiative thanks to the Copenhagen Consensus Center; and the FHI has advised several organizations at all levels, from the U.S. State Department and the UK Prime Minister's Office to the United Nations and the World Economic Forum on X risks (Duda, 2018).

As Duda (2018) explained, GiveWell's work in identifying outstanding global health initiatives and bringing the effects of policies for future generations into consideration are all examples of how small amounts of global priority research can have an immense impact. As a relatively new field, effective altruists expected to find many more low hanging fruits for global priorities research (Duda, 2018). If politicians and decision-makers are presented with good arguments and enough evidence, effective altruists expect to be able to steer these funds into effective causes and bring important issues to the agenda.

The only major academic center working on this is the Global Priorities Institute at Oxford. Research projects in this field are also conducted by non-academic organizations such as the Open Philanthropy Project and in part at 80,000 Hours (with a stronger focus on career strategy). But there are also independent researchers, such as Carl Schulman, who perform freelance work for donors and organizations pursuing global priorities. The work of the Future of Humanity Institute on macro-strategy research in long-term outcomes of present-day actions is also relevant for this field. Other organizations relevant but not directly working on this cause are the Cambridge Center for the Study of Existential Risk, analyzing different risks to humanity, and the Copenhagen Consensus Center. The last of which occupies itself with assessing solutions to global problems, with particular interest on international development. The Global Catastrophic Risk Institute, Foundational Research Institute and the Future of Life Institute, all working on X risks are also considered relevant for this cause area (Duda, 2018).

6.5 Non-human animals

The question is not, can they reason? Nor, can they talk? But, can they suffer?
(Bentham, 1789, p. 144)

The industrial revolution has brought many things with it, from machinery and chemical manufacturing to today's automatized assembly lines. This rationalization process has affected all parts of our lives, even the ways we produce our food. The agricultural has not only massively reduced the prices of crops, but also of animal products. Today, billions of animals are raised every single year only for our consumption. In 2019, 69 billion chickens, 1.5 billion pigs, more than 655 million turkeys, 564 million sheep, 470 million goats and 298 million cows were raised, killed and sold as meat (Our World in Data, 2018). In most cases the animals live and die under horrendous conditions in intensive animal farming facilities.

As for the number of fishes slaughtered every year, there are no exact numbers as they are usually counted in tonnages and not individually. But the most recent estimates, using the estimated mean weight of the fishes by the tons caught, account for a total estimate of 0.79 – 2.3 trillion wild fish caught for human consumption in 2019 (fishcount.org, 2019b). The yearly estimates of fish raised in aquaculture are between 47.8 – 158.7 billion fish (fishcount.org, 2019a).

Whether the suffering of non-human animals even matter is a debate is as old as history. Already in Ancient Greece did Pythagoras (c. 580-c. 500 BCE) speak against the killing and eating of animals (Ann Violin, 1990, p. 123). Certainly, there have been many contemporary scholars and thinkers arguing for the inclusion of non-human animals into our moral circle for years. Maybe the most prominent one has been Peter Singer with his book *Animal Liberation*, published in 1975. The arguments are already well-known, but I will shortly summarize the main reasoning and expand on the latest moral circle expansion animal ethicists are currently discussing.

6.5.1 Do non-human animals count?

There is a wide consensus in academia that at least higher animals are able to consciously experience affective states, such as suffering. The Cambridge Declaration on Consciousness published in July 2012 is just one example of this agreement on non-human animal consciousness in the scientific community:

We declare the following: “The absence of a neocortex does not appear to preclude an organism from experiencing affective states. Convergent evidence indicates that non-human animals have the neuroanatomical, neurochemical, and neurophysiological substrates of conscious states along with the capacity to exhibit intentional behaviors. Consequently, the weight of evidence indicates that humans are not unique in possessing the neurological substrates that generate consciousness. Non-human animals, including all mammals and birds, and many other creatures, including octopuses, also possess these neurological substrates.” (Low, 2012, p. 2)

Clearly, the scientific approach always entails a degree of skepticism to all scientific conclusions, including this one. However, this healthy skepticism should certainly not be overinterpreted, since it is safe to say that most scientists agree that non-human animals have conscious experiences. As a matter of fact, we have not even been able to demonstrate human consciousness with our current set of tools, but you would most certainly not question that your neighbor lacks consciousness but rather act as if it were a certainty (Tomasik, 2017). Several philosophers argue that when it comes to non-human animals, it is a similar case. And as Brian Tomasik (2017) argued, even with odds of around 80-85 percent of certainty that non-human animals having conscious experiences, in practical terms it is better (from an ethical standpoint) to assume that they have consciousness.

As an analogy, we only need to look at the debate on climate change to see how the communication to the public of a healthy level of skepticism can backfire at a social level (Tomasik, 2017).

Now, with this in mind, there might still be people who think that even if non-human animals can suffer, they may suffer less intensely than humans. To this, Tomasik (2017) responded that, it is unclear whether animals really do have substantially less sophisticated consciousness than humans. They indeed show many of the psychopathologies we humans do and are also used as model organisms for depression in drug trials. We now know that some animals, such as elephants and chimpanzees have well developed cultures that differentiate between groups (Balter, 2013).

Even if it were the case that non-human animals suffer less, let us say by one or two orders of magnitude, the numbers of individuals suffering at the hands of humans are just too vast to ignore. Some, such as Tomasik (2017), even argue that it might be the case that non-human animals suffer even more than humans. As they lack those higher-level thoughts, we humans might have during moments of pain, they may experience pain in absolutely raw terms.

6.5.2 Especially neglected cause

Due to the immense number of animals involved, any type of progress in this area could avoid a huge amount of suffering. The problem is not only of massive scale, but it is also extremely neglected. In 2017, only 3 percent of all donations in the U.S. went to charities that focused on the environment and animals (Charity Navigator, 2018). And from this, merely 1 percent went towards farmed animals (Animal Charity Evaluators, 2016). Figure 10 illustrates the stark contrast between the numbers of animals suffering and the donations towards the charities helping those animals in the U.S for that year.



Figure 10 Number of animals used and killed by humans in contrast to the money donated to animal charities in the U.S.. Figure adopted from Animal Charity Evaluators (2016)

While “[...] over 99.6 percent of all domesticated land animals used and killed by humans, 0.2 percent are animals used in laboratories, 0.07 percent are used for clothing, and 0.03 percent are killed in companion animal shelters” (Animal Charity Evaluators, 2016). However, most donations to animal charities (66 percent) are for companion animal shelters and 32 percent are for organizations doing mixed or other activities. Just 0.7 percent go to organizations fighting for laboratory animals, and only 0.8 percent go specifically to organizations working to help farmed animal (Animal Charity Evaluators, 2016).

6.5.3 Wild-animal suffering

The most recent dilemma in animal ethics in the Effective Altruism community entails the question whether we should also care about the suffering of wild animals. Whether an animal is born in either a domestic or in a wild environment is completely random, and we would certainly not care less about the wellbeing of a cat or a rabbit just because it was born in the forest. Now, why should we care less about the suffering of deer or foxes, just because they were born in the wild? In fact, people were devastated by the shocking pictures of wild animals suffering and dying during the wildfires in Australia from July 2019 to March 2020. Causes collecting funds to save them and help reforest their habitat raised thousands of dollars. The WWF alone received more than eight million dollars for their recovery programs (WWF, n.d.).

Just as philosophers argue that treating non-human animals worse than other humans is based on a prejudice based on species – *speciesism*, – the line between *pets* and *farmed animals* but also *domesticated* and *wild* animals seems unjustified (Dunayer, 2004; Singer, 2009; Vinding, 2015).

When taking wild animals into account, the urgency to do something to alleviate the suffering of non-human animals becomes even more imperative. They are by several orders of magnitude more numerous than non-human animals on factory farms, in laboratories, zoos, and those kept as pets combined (Tomasik, 2009). The exact numbers of animals living in the wild remain unknown, but Tomasik (2009) provided us with some rough estimates. According to his calculations there are between 10^{11} and 10^{14} land vertebrates and at least 10^{13} marine vertebrates. These numbers are only for vertebrae and exclude protostome (mollusks, arthropods, and annelids) and thus animals such as octopi, crabs, insects and arachnids, whose levels of consciousness remain a debate in the scientific community. Even excluding them, the numbers of individual animals in the wild is enormous and the problem is that:

Most of these animals endure intense suffering during their lives, such as from disease, hunger, cold, injury, and chronic fear of predators. Many wild animals give birth to tens or hundreds of offspring at a time, most of which die young, often in painful ways. This suggests that suffering plausibly dominates happiness in nature. (Tomasik, 2015, p. 133)

As discussed above, the cause to help farmed animals, animals who are suffering on our behalf, is extremely neglected. But, as we might imagine, wild-animal suffering is an even more neglected problem. Not only is the academic discussion very recent and the size of the problem enormous, but we also lack the knowledge and understanding necessary to know how we could best help animals living in the wild. When it comes to wild animals, most research has primarily focused on either its preservation or its exploitation for human interests. The latter approach sees nature and the animals living in it as resources whose use and consumption should be coordinated among people. While those arguing for nature's preservation would prefer human intervention to be reduced to an absolute minimum and let nature "take its course". Neither of these dominant approaches addresses the animal welfare perspective at all and while the discussion of animals' rights and interests had already entered the field thanks to Peter Singer (1975), it was never intensely discussed (Gamborg et al., 2012).

6.5.4 Organizations working on their behalf

Similar to GiveWell, the meta-charity Animal Charity Evaluators (ACE) evaluates the work of organizations working to reduce the suffering of non-human animals. They also publish a yearly report on the most effective organizations according to their criteria, which are similar to GiveWell's. Their latest four top charities for 2020 entailed two organizations working for farmed animals, the Albert Schweitzer Foundation and the Humane League. Both of which are dedicated to improving the welfare standards of farmed animals in the agricultural industry via corporate outreach, advocacy, and legal work.

ACE also recommend the Good Food Institute as one of their top charities. This organization is dedicated to transform the animal agriculture industry by accelerating the development of alternative proteins and are active all around the world, connecting and supporting research in plant-based protein alternatives as well as cellular culture, including clean meat. The latter would be especially revolutionary as it would completely avoid the necessity to give up animal products altogether as they would be completely replaced one by one, only excluding animals from the equation altogether.

The fourth organization they promote is Wild Animal Initiative, which promotes the establishment of an academic field for wild animal welfare, which as discussed above is an especially neglected topic in animal ethics, wildlife management and conservation. "[Their] mission is to understand and improve the lives of wild animals" (Wild Animal Initiative, n.d.).

Wild Animal Initiative tries to encourage research and public interest in this field for the development of solutions to the problems wild animals face (Wild Animal Initiative, n.d.). Their latest report, for example, discussed the idea of wildlife contraception as a more humane alternative for population control (Wild-Animal Suffering Research, n.d.).

6.6 Other cause areas

The four cause areas highlighted above are certainly the cause areas most effective altruists (at least those in the latest EA survey in 2019) care about. There are, however, many more causes that despite not being the main focus, other in the community are currently working on. These include according to the Centre for Effective Altruism (EffectiveAltruism.org, 2020):

- Improvements to the scientific establishment, such as greater transparency and replication of results
- Researching mental health and neurological disorders, particularly depression and anxiety, and improving access to treatment in low-income countries
- Tobacco control
- US criminal justice reform
- International migration and trade policy reform
- Electoral reform (e.g. implementing better voting methods)

6.7 In the classroom

As the time available for students at school to discuss complex and expansive topics as Effective Altruism, it is necessary to leave some aspects of EA out of the classroom. Similar to the cause area prioritization effective altruists perform, the aspects that could potentially be discussed in the classroom need to be prioritized as well. Personally, I think that the subject of cause areas in EA should be discussed in class, but not as extensively as the subtopics discussed above. Thus, it might be best to separate the class into five groups and ask each of them to elaborate and present one cause area shortly presented by the teacher in forehand. They could be encouraged to estimate which cause areas receive the most attention by the EA community and obtain points for correctly ranging the first top five cause areas among the top five (see Figure 5). If there is enough time left after their presentations, they could also discuss whether they themselves agree with the prioritization of the EA community. Chapter 8 entails a lesson plan for this exact activity.

Further thoughts:

The cause area on non-human animals can be easily linked with the topics, non-human animal ethics and vegetarianism/veganism. However, as non-human animal ethics should not be discussed merely in the periphery, I would discuss the cause areas only at the end – after having discussed the criticism of Effective Altruism. This would allow for a seamless transition from non-human animal welfare into animal ethics.

As discussed above, even effective altruists cannot fully agree on the prioritization of cause areas. This is why, some of them argue to distribute their donations to more than just one cause area to reduce their clustering risk. This risk diversification approach can also be found in financial investment, e.g., investing in exchange traded funds (ETFs) instead of one single stock; an approach that will certainly come in handy to students, when it comes to other life areas, such as private investments.

7 Criticism of Effective Altruism

It was to be expected that a social movement criticizing and re-evaluating the way we do charity, would be met with a lot of backlash. Effective Altruism is seen by some social justice activists as insensitive to justice and discriminating against the worst-off. Others have a problem with the normative and methodologies EA applies to evaluate charities. There is also the worry of the systemic impact Effective Altruism might end up having, shifting resources to those organizations it considers effective (Gabriel, 2017, p. 458). But certainly, the strongest critique is that EA is only addressing the consequences of a failed system and fails to address the main problem: the system (Berkey, 2018, p. 143).

There is also an article in the *Stanford Social Innovation Review* blog from Ken Berger, the CEO of Charity Navigator and his colleague Robert M. Penna throwing several very impetuously formulated reproaches at Effective Altruism. These include some of the aspects discussed by Gabriel (2017) but their strongest criticism lies in the “elitist” nature of the movement, which according to them might end up having seriously negative implications in the charity sector (Berger & Penna, 2013).

7.1 Implicit assumptions within Effective Altruism

The core message of Effective Altruism is that “we should do the most good we can” with our spare resources to make the world a better place. This premise as such is compatible with various moral theories and does not yet give any specifics about the type of help nor the individual’s relationship to it (Gabriel, 2017, p. 458).

But Effective Altruism, as elaborated in this thesis, does not limit itself to this core message. The movement is as MacAskill (2019) argued, based around two projects, one intellectual and one practical, i.e., the implementation of the core idea. And it is here, where many authors argue that a number of different assumptions are presented as implicit. Just taking a look at Peter Singer’s *The Life You Can Save* (2010) and William MacAskill’s *Doing Good Better* (2016), there are some evident premises I elaborated on in this thesis.

Gabriel (2017) summarizes these premises of the “thick version” (as he called it) of Effective Altruism as follows:

First, it adopts a largely welfarist understanding of value. According to this view, good states of affairs are those in which suffering is reduced and premature loss of life averted. Second, it is broadly consequentialist, maintaining that we should do whatever maximizes the sum of individuals welfare at all times. [...] Third, the movement takes ‘a scientific approach to doing good’, which means using tools such as cost-effectiveness analysis and randomization to help quantify and compare the impact of different interventions. (Gabriel, 2017, p. 458f)

7.1.1 The dominance of consequentialism

Just taking a look at what Holden Karnofsky, GiveWell’s co-founder writes about Effective Altruism is a good illustration of these assumptions. For him Effective Altruism “[...] is a way of thinking about morality that insists on the *maximization* of good accomplished and not just *satisficing* of rules and guidelines” (Karnofsky, 2016). And indeed, according to the statistics of the last EA survey (2019) named above, the consequentialist tendency in the movement becomes evident. The majority (80.7 per cent) of the 1,360 respondents identified with consequentialism, especially utilitarian consequentialism, while virtue ethics and deontology made up only 7.3 and 3.2 per cent respectively (Dullaghan, 2019).

Consequentialism regards that normative properties of a given action (or inaction) depends only on its consequences. This school of thought is not only historically important, but it remains one of the most prominent influences in the field of morality and the rightness of acts. It regards only the consequences as relevant for the evaluation of the act and not its motives or its alignment to a general set of rules. Utilitarian consequentialism, classically represented by Jeremy Bentham (1789), John Stuart Mill (1861), and Henry Sidgwick (1907) (Sinnott-Armstrong, 2019). Today’s most prominent utilitarian is Peter Singer, active in all realms of ethics from animal rights to poverty alleviation, he could be considered as one of the most influential contemporary philosophers.

In part due to its historical importance, consequentialism has also had many critics over the years. One of the most concise criticism is probably that from Rudolph (2008), who highlighted the problems with consequentialism that arise from its lack of boundaries in time and scope of effects.

Consequentialism fails to define the point in time from which an action is to be evaluated. There is also no constraint in the scope of effects following a given action to be considered for the judgement of an action (Rudolph, 2008, p. 67). To illustrate this gap in consequentialist thinking he used the Baby Hitler scenario:

It is a bright summer day in 1895 and you are walking across the bridge that spans the river in Linz, Germany. You notice some children playing near the mill and take a moment to watch them romp by the water's edge. Just then, one of the smaller boys slips into the water and is taken away by the current. It is clear that he cannot swim and will drown if you do not save him. You jump in the rushing water [...] and swim him to safety. The next day, you are a city hero, as you have saved the six-year-old son of Klara Hitler, Adolph. (Rudolph, 2008, p. 68)

Now, is the action of saving Baby Hitler's life a good one or a bad one? Clearly, at the time of the rescue it would be considered a good action. However, the consequences of that action radiate outwards from that day onward. However, as time went on, which consequences of the rescue act are you responsible for? Is it limited to the consequences in the few days that followed the event, or does it include his actions in the years after that or even those until his death in 1945? Without these criteria in time and scope, Rudolph (2008, p. 72) argued that consequentialism cannot stand and is thus incapable of providing ethical imperatives. William H. Shaw's defense of consequentialism establishes these limitations in a somewhat intuitive way:

[...] at the end of the individual's foresight [...] By limiting an individual's responsibility at the extent to which the actor could have predicted the consequences of an act, they can only be held responsible for that which they could have made an informed decision. (Rudolph, 2008, p. 72)

However, Rudolph (2008, p. 72) identified a problem with this logic. If the evaluation of an individual's actions depends on how much they know when making the decision, they have an incentive to make decisions as uninformed as possible. And this would probably have negative consequences consequentialists would prefer to avoid.

Another commonly highlighted shortcoming of consequentialism comes from the deontological ethics, which criticizes consequentialists for ignoring relevant aspects on their strive for the maximization of good. Deontologists value the nature and quality of relationships among people as well as the distribution of good, figures they regard as inherently valuable unlike consequentialists. Heard (1997), for example, argued that human rights are the best way to protect minorities and ensure that no-one is subject to any type of utility calculations for the “greater good”. “If anything is to stand in the way of governments or societies sacrificing individual or minority interests in favour of the collective, it is the bulwark of human rights” (Heard, 1997).

And on a much deeper level, Gabriel (2017, p. 461f) highlighted the failure of Effective Altruism to take the possibility of moral error into consideration, as “[...] consequentialism may not be the correct account of personal morality, so conclusions that assume its truth are worryingly contingent” (Gabriel, 2017, p. 462). This also might cause the unintentional alienation of potential supporters who share the core message of Effective Altruism of “doing the most good you can” but endorse other ethical values, such as fairness and equality, that are evidently underrepresented in the movement (Gabriel, 2017, p. 462).

Getting into detail about the debate on the moral error theory would certainly go beyond the scope of this thesis. But in brief, Mackie’s argument for the moral error theory is based on the empirical observation that there is huge variety in moral views around the world, which “[...] reflect adherence to and participation in different ways of life” (1977, p.36). Due to this, he argued, it is impossible to determine which moral theory has the “facts” about morality and would thus be considered to be “objectively valid”, as each view on a given topic develops in a cultural context but not in the other (Mackie, 1977, p. 29f).

7.1.2 Welfarism

In connection with the consequentialist tendency of Effective Altruism, the *welfarist understanding* Gabriel (2017, p. 458f) highlighted is evident in many publications and organizations around Effective Altruism. William MacAskill's *Doing Good Better* (2016) elaborated above is a good example for this. So is Peter Singer's 2010 *The life You Can Save* and 2015 *The Most Good You Can Do* and *Famine, Affluence, and Morality*. So, do Magnus Vinding *Effective Altruism: How Can We Best Help Others?* (2018).

Even though neither of the authors named above name it explicitly, the focus on welfare is omnipresent and it is safe to say that it is indeed one of the basic building blocks of the whole movement. As Nick Gabrieli (2021) wrote: "One of the key philosophical tenets of the Effective Altruist movement is "welfarism," the moral theory that the well-being of people is the sole or at least a major component of what is morally desirable." However, welfarism and its main economic assumptions of utility can be regarded critically as well, as Brayan Viegas Seixas (2017) put it:

[W]elfarism has three main problems: (1) its central idea of overall sum of individual utilities does not capture societal values decisively relevant in the context of health; (2) the use of the Potential Pareto Improvement brings an unresolvable separation between efficiency and equity; and (3) individual utility may not be a good measure in the health sector, given that individuals might value things that diminish their overall health. (Seixas, 2017, p. 1)

The dominance of the welfarist approach in health economics overshadows the multiple possible alternatives that are still being discussed in academia. From several concepts under the Multi-Criteria Decision Analysis, which take a wide variety of outcomes and principles into consideration, such as communitarian values of groups or societies (Seixas, 2017, p. 7).

7.1.3 The Scientific Approach

Ultimately, wanting to assess the size of the problems as well as the impact of the interventions to solve them, effective altruists embrace the scientific method. So, yes, the scientific approach as elaborated in the chapters above on Effective Altruism is also one of the movement's basic building blocks (MacAskill, 2019a).

But this rational scientific approach has certainly met critics as well, as not everyone agrees that charity should be evaluated and rated as Effective Altruism does. Martin Morse Wooster (2017) from Capital Research Center argued:

[T]he problem is that there is no one best way to give and giving is an art, not a science that can be defined by firm and inflexible rules and derived from axioms everyone accepts. Pretending giving is a science does not make it a science. I also question the calculus that the only way to give is to give the greatest good to the greatest number. Effective Altruism seems to me to be a cold and joyless way to give. It denies the giver the personal joy that comes when they help another person that they know. (Wooster, 2017)

Ken Berger, CEO of Charity Navigator and his colleague Robert M. Penna put similar thoughts in a more brute manner in their critic of Effective Altruism in the *Stanford Social Innovation Review* blog in 2013. They argued that comparing causes against each other “[...] amounts to little more than charitable imperialism, whereby “my cause” is just, and yours is—to one degree or another—a waste of precious resources” (Berger & Penna, 2013). In correspondence with William MacAskill, they criticized Effective Altruism as immoral for trying to find the most effective causes as they hold the view that “[...] it's impossible to weigh one person's interest against another's [...]" (MacAskill, 2016, p. 40). In their piece, Berger and Penna (2013) appropriated Effective Altruism as “defective altruism” and criticized:

[...] defective altruism is — by the admission of its proponents — an approach that not only unjustifiably claims the moral high ground in giving decisions, but also implements this bold claim by weighing causes and beneficiaries against one another. In this, it is not moral, but rather, moralistic in the worst sense of the word. (Berger & Penna, 2013)

It is maybe important to note that Charity Navigator, founded in 2001, is a charity dedicated to assessing and evaluating thousands of organizations based in the U.S. They promote their material with slogans such as “give smarter” (Charity Navigator, 2021) and argue to have

developed an “*unbiased, objective, numbers-based rating system* to assess over 160,000 of America's [...] charities” (Charity Navigator, n.d. italics added for emphasis). They create these rankings by rating two areas of a charity’s performance: “[...] their Financial Health and their Accountability & Transparency” (Charity Navigator, n.d.). Their ratings “[...] show givers how efficiently [Charity Navigators] believe a charity will use their support today, how well it has sustained its programs and services over time and their level of commitment to good governance, best practices and openness with information” (Charity Navigator, n.d.).

Clearly, there might also be people displeased by the thought of putting a monetary value on someone’s life. GiveWell, for example, estimates the costs of saving one life in the developed world, through donations to the charities it recommends to be approximately \$2,000 (Hassenfeld, 2007). Pete Mills questioned the preciseness of Effective Altruism’s calculations (Todd et al., 2021, p.6) on how much it costs to save a life.

7.2 The Institutional Critique of Effective Altruism

By far the most common objection to Effective Altruism has been its lack of systemic critique (Berkey, 2018). When it comes to poverty alleviation, many see its efforts in promoting mainly organizations helping the poor directly as an approach merely combating the symptoms of a much bigger problem. A problem that according to them, Effective Altruism fails to address (Berkey, 2018, p. 149). As Brian Berkey (2018) summarized in his review of *The Institutional Critique of Effective Altruism*:

This complaints [Effective Altruism’s failure to address institutional change] has been made in somewhat different ways by different critics, but the unifying thought seems to be that there is something about institutional change that is distinctively morally important in the face of the facts on the ground regarding global poverty, and that this is something that effective altruists have tended to overlook, at least in part due to their acceptance of certain moral and/or empirical claims that ought to be rejected. (Berkey, 2018, p. 145)

Lisa Herzog, postdoctoral researcher at Goethe University in Frankfurt, argued that “[...] the most dangerous underlying assumption in the worldview of effective altruists [is] that they take the current institutional order as given, implicitly denying that it is open to change” (Herzog, 2016).

According to Judith Lichtenberg, professor of philosophy at Georgetown University, Effective Altruism fails to recognize the value systemic change has in improving people's lives, especially for the poorest individuals:

[T]he maximum effectiveness strategy means neglecting programs that support advocacy for political and structural change, which are essential for addressing the deeper roots of poverty. [...] People across the political spectrum should agree that structural changes that allow all workers to earn a decent living are preferable to welfare programs and private charity. (Lichtenberg, 2015)

Similarly, Amian Srinivasan argued in his critic of MacAskill's *Doing Good Better* is problematically complacent when it comes to the institutions of global capitalism. She recognized that some people around Effective Altruism have taken the promotion of institutional change into consideration when thinking about ways to make the world a better place. Examples are organizations taking up on the reform of immigration laws, penalty law and factory farming. But for her, it is troubling that "[e]ven in these cases, the numbers are what matters" (Srinivasan as cited by Berkey, 2018, p. 151). Herzog and Srinivasan questioned the expected value calculations MacAskill (2016) did for evaluating the impact of potential career paths. In the case elaborated in his book, he was able to put some numbers on his estimations, but what about evaluating more radical political change practices:

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 unrecognizably different world – even, perhaps, on our becoming unrecognizably 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 reorganizing society? (Srinivasan as cited by Berkey, 2018, p. 152)

According to Srinivasan, Effective Altruism's focus on finding the most effective and efficient ways to do good in itself is capitalistic in nature. She argued that MacAskill's uses of expressions typical from global capitalism and his focus on quantifying and productivity, to be the result of an intrinsic relationship between the movement and capitalism (Srinivasan as cited by Berkey, 2018, p. 152). "That he speaks in the proprietary language of the illness – global inequality – whose symptoms he proposes to mop up is an irony on which he doesn't comment" (Srinivasan as cited by Berkey, 2018, p. 152).

She suspected that reasons for Effective Altruism, but more specifically for MacAskill to not address this issue might lie with the nature of his “[...] potential followers – privileged, ambitious millennials – [who] don’t want to hear about the iniquities of the system that has shaped their worldview” (Srinivasan as cited by Berkey, 2018, p. 152f).

Furthermore, Srinivasan argued that Effective Altruism is a conservative movement, embracing the status quo. She further stated that:

MacAskill does not address the deep sources of global misery...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 as cited by Berkey, 2018, p. 153).

7.2.1 The problem with earning to give

In line with the criticism on Effective Altruism negligence of systemic change, many scholars regard earning to give (also called professional philanthropy) as similarly problematic. They criticize Effective Altruists as “[...] unduly complacent toward the prevailing institutions of global capitalism, or more strongly, have in fact allied themselves with those institutions, despite their role in perpetuating the suffering that they aim to relieve.” (Berkey, 2018, p. 150). They see the concept of earning to give as problematic, as it recommends people to choose high-earning careers, such as investment banking, to be able to donate even more to effective charities. Peter Mills argued in a discussion about 80,000 Hours with Benjamin Todd, Sebastian Farquhar and Tom Cutterham in *The Oxford Left Review* in 2012 that:

[...] professional philanthropy [that is, earning to give] does not just involve making your peace with the system – it means embracing it. The unstated imperative: don’t rock the boat [...] what enriches you is your position in a set of profoundly exploitative social relations, which we might label capitalism. (Todd et al., 2012, p. 5)

By linking the impact of the *professional philanthropist* (someone doing earning to give) with their salary, the individual’s interests are consequently aligned with the company’s interests. So, anything that might affect your company’s income negatively, such as taxes or regulations might become an issue. Pete Mills argued that this only reinforces the problems of the system we live in. “The good that a professional philanthropist does depends on perpetuating a system which immiserates a vast portion of the world’s population” (Todd et al., 2021, p.5).

There are also critics, such as Berger & Penna (2013) from Charity Navigators, that take the idea of earning to give as absurdum. According to them, discouraging people from taking up jobs in the non-profit sector might end up leaving even those organizations Effective Altruism supports without staff (Berger & Penna, 2013).

7.2.2 Counterargument to the institutional critique

Berkey (2018) denoted the systemic criticism of Effective Altruism as somewhat shortsighted. Certainly, the most prominent topics in Effective Altruism, especially those that are most strongly used in outreach, have a strong focus on charities working to help the world's poorest people. But, as he argued (2018, p. 162f) this focus on evidence-based programs, especially on RCTs, is mainly to counteract the common belief among people that donating to charity does not change anything. Effective Altruism is in fact, much more than just the Against Malaria Foundation and the Deworm the World Initiative.

Taking a deeper dive into Effective Altruism allows us to recognize that the topics it deals with are much broader than just helping the poor. Organizations in the Effective Altruism movement, such as the Open Philanthropy Project, attempt to find solutions to a range of policy areas. The Future of Humanity Institute and the Future of Life Institute focus on existential risks, such as many others discussed in chapter 5. Berkey (2018, p. 155) asserted that there is nothing in Effective Altruism that might hinder people in the community to endorse efforts to promote institutional change motivated by good moral reasons. In fact, assessing the expected value of enrolling in a political career or becoming an activist is exactly what effective altruists recommend us to do to find new effective ways of improving the world (Berkey, 2018, p. 155).

Similarly, Benjamin Todd and Sebastian Farquhar responded that there is no such thing as “one dominant solution” (Todd et al., 2012, p. 5) for choosing a career path such as those proposed by 80,000 Hours. They defended the notion of earning to give as being a good way to do more good than working directly in charity only in some cases but not always (Todd et al., 2012, p. 5f), as it is also presented in 80,000 Hours. Earning to give is not a one-fits-all solution, but more a reminder that there are other ways to help than working directly at NGOs.

7.3 Elitist Philanthropy

The last of the most common criticisms of Effective Altruism is its branding as an elitist movement. Having its origin at the University of Oxford and being represented overwhelmingly by people with academic backgrounds – 90 percent have attained or are about to complete a post-secondary degree and 20 percent have attended one of the top 20 universities in the world for their undergraduate studies (Dullaghan, 2019) – this critic is understandable.

Maybe the noisiest critic of Effective Altruism in this regard, stems from Berger and Penna (2013), who work at Charity Navigators. They fear Effective Altruism to start a tendency towards “[...] a more centralized form of giving where experts decide where the money goes, rather than individual donors” (Berger & Penna, 2013). They see EA’s reliance on data and science as elitist and even arrogant. For this they highlighted Eric Friedman’s 2013 book passage: “Though not *necessarily* morally superior to do-gooders, do-besters may be *intellectually* superior [...]” (Friedman, 2013 as cited by Berger & Penna, 2013).

Berger and Penna (2013) also fear for “[...] volunteerism for the general public, since most “worthy” efforts are overseas and require a professional degree to have what Friedman calls “deep expertise in niche areas”?” (Berger & Penna, 2013).

This point of critic is a common dilemma discussed between effective altruists since the beginnings of Effective Altruism. Some question the efforts of others to spread the word and motivate as many people as possible to “do the most good they can”. They argue that due to the unequal distribution of people’s range of possibilities for action – that is their wealth, abilities, and power – the movement should “[...] prioritize outreach for primarily the most effective people” (Diogenes, 2016), because:

If your goal is to have as much money as possible be donated in the best ways possible, then you should seriously consider whether the expected value of recruiting Bill Gates or other billionaires is higher than the expected value of recruiting as many people as possible. (Diogenes, 2016)

Similar to wealth and most commonly related to it, *power* – the ability to significantly influence others’ actions – is unequally distributed in society. So, according to Diogenes (2016) effective altruists should focus on reaching out the most influential people to maximize their impact. He recognized that this might sound counterintuitive or even “inhumane” to some, but he argued that “[...] an EA [effective altruist] chooses to prioritize for the most humane reason possible, almost by definition” (Diogenes, 2016).

At the same time, other effective altruists oppose this position, as they recognize the damage it might cause to the movement as a whole. Not only due to the negative public perception it could generate but it also might chase people away that would otherwise have joined the movement and add significant value to it. This might also include people from minority groups, whose perspectives are certainly underrepresented in Effective Altruism as a whole. On a less consequentialist line promoting diversity and fairness might just be “good to do”, especially so, if the decisions that are being taken, might in some form or another affect the individuals included at a personal level (Diogenes, 2016).

7.4 In the classroom

To discuss the criticism of Effective Altruism in ethics class is certainly the most challenging part of the lesson series. Most of the critic points demand prior knowledge in several areas, e.g., capitalism critic, ethical theories, welfare economics, etc. A comprehensive discussion on each point of critique point discussed above might very likely extend beyond several teaching units; and the list presented above is certainly not exhaustive. There is certainly enough material discussing the criticism of EA for a whole master’s thesis to be written on the topic.

There are several ways to deal with this problem, which are necessarily dependent on how much time one would like (and is able) to dedicate to the topic. Personally, I would discuss the points of critique short and concisely, and only elaborate on one of the points, if students seem particularly interested in the topic. Should the interest be bigger, I would propose the pupils to prepare a presentation on the point of critique to discuss it with their classmates in one of the lessons. This presentation would certainly account as substitutive performance for other mandatory performances in class.

8 Effective Altruism in ethics class

8.1 The Austrian ethics class curriculum

At the moment of writing, the Federal Ministry for Education, Science and Research is still working on the ethics class curriculum for the next schoolyear 2021/22 (BMBWF, 2020). However, a curriculum draft is already available on their website, to which the subject of Effective Altruism seems to fit perfectly. Among the already defined goals for ethics class are (BMBWF, 2020, p. 1):

- Valid argumentation and reflection of ethics and moral questions
- Empower students to autonomous reflection in regard to successfully shaping their life for the better and guide them through fundamental questions on their own existence and society in general
- Enable students to analyze their personal experiences and encourage them to critically self-reflect their judgements and actions
- Boost the students' disposition to critically examine their own attitudes and moral (pre-)judgements

Furthermore, discourse orientation is among the didactic principles mentioned in the curriculum draft: "Possible solutions for moral and lifestyle problems are to be discussed discursively, more specifically, suggested responses are to be critically analyzed. For this end, diverse formats of conversation and discussion are to be made use of" (BMBWF, 2020, p. 3 transpalted by the author).

As illustrated in the elaboration on Effective Altruism in the chapters above, including the pedagogic comments, the subject of EA seems to be especially suitable for ethics class. It not only fits the ethics class curriculum's goals named above but it also brings a wider, global perspective into the classroom.

As evident as it might seem to teachers outside of the classroom, it is especially important to keep the indoctrination prohibition in mind during class as well. "Every type of indoctrination, no matter how subtle, is to be refrained from in ethics class. The judgement formation of the students shall not be influenced by particular accentuations of any worldview nor religious standpoint of any kind on the part of the teacher" (BMBWF, 2020, p. 3 transpalted by the author).

This master's thesis shall not constitute neither an apologia nor publicity for Effective Altruism. Similarly, the indoctrination prohibition is certainly to be respected when discussing the subject of Effective Altruism in ethics class.

8.2 Lesson plans

Preliminary remarks:

- The complexity of the subject requires a specific level of maturity, education level and constructive discussion culture from the students. Therefore, I would recommend discussing this topic in the 11th or 12th school grade (equivalent to 12th and 13th grades in higher-level technical and vocational colleges [ge. Berufsbildende höhere Schulen]).
- For the second lesson (the logical argument for Effective Altruism) a minimal understanding of logic and argumentation theory are necessary to foster valuable discussions.
- Lesson units in Austria have a duration of 50 minutes each. However, from personal experience, the actual duration of a lesson is less than that. For this reason, the lesson plans elaborated below are planned for a duration time of 45 minutes.
- Abbreviations for the lesson plans: S = students; T = teacher; PPP = Power Point Presentation

8.2.1 Lesson plan 1: The ethical argument for EA

time (min)	phase	content	arrangement	media
5	guidance	T presents the subject (EA) to be discussed during the next lesson units and presents either one of its definitions (see chapter 3) or asks the students to search for the term in a search engine and read 2-3 definitions out loud (Wikipedia, CEA, etc.).	Discourse by T & eventually individual work + discourse by S	For the presentation of the definition of EA, T can use either a PPP or the board and chalk. If S are asked to search for the terms, they will need access to the web via web-enabled devices, such as a smartphone.
5	guidance	Peter Singer's thought experiment (child in the pond) is presented in video form on a projector.	Watching the video together	Video (https://youtu.be/7bwM_atwgye) & projector
25	elaboration	Discussion of the thought experiment in plenum. Optional: present and discuss further thought experiments discussed in chapter 4.	Discussion in plenum (T as moderator, participates in the discussion and corrects if necessary)	
10	securing knowledge	2-3 S summarize the results of the discussion; T adds comments if necessary. Everyone takes notes.	Presentation by T & eventually discourse by S + individual work	Writing materials and or notebook

8.2.2 Lesson plan 2: The logical argument for EA

time (min)	phase	content	arrangement	media
3	guidance	1-2 S repeat the results of the last lesson; T adds further comments if necessary.	Discourse by S & eventual discourse by T	Notes made by S
5	guidance	T writes Peter Singer's logical argument (chapter 4) on the board; S make notes.	Discourse by T	Board, chalk Tafel, and writing material and or notebook
33	elaboration	method: Pair, Square, Share S first discuss the logical argument with the person sitting next to them (5 min), then with the pair in front / behind them (8 min); results are then shared with the whole class to start the discussion in plenum (20 min).	Working in pairs & groups & plenum discussion (T as moderator, participates in the discussion and corrects if necessary)	
4	securing knowledge	S write the most important arguments down & receive the homework to discuss the logical argument with their parents.	Individual work: S discuss the topic with their parents	Writing material and or notebook

8.2.3 Lesson plan 3: Among the 10% & Five key aspects of Effective Altruism

time (min)	phase	content	arrangement	media
5	guidance	1-2 S repeat the results of the last lessons, T comments if necessary & 1-2 S share with the class about their discussion with their parents (homework)	Discourse by S & eventual discourse by T	Notes by S
15	elaboration & securing knowledge	T draws the X-axis from Figure 1 (Global Income Distribution in 2011) on the board and asks S, where they would estimate their current and future households to be. The results are then discussed with the actual figure, and the potential ethical implications of the income distribution are discussed. Optionally: T can make the claim that it should be the richest 10 percent of the world population, who should be held accountable to help those in need, and not necessarily the average citizens (see comment in segment “Among the 10%”).	S-T-discussion	Figure 1 (Global Income Distribution in 2011) & board and chalk
1	guidance	T continues with the <i>Five key aspects of Effective Altruism</i> by saying „Should you decide to donate a share of your future income and or volunteer for a project working on making the world a better place, the following questions might help you decide where to start: “Where are these resources currently most needed? What should be prioritized and how could we help? These are the leading questions of the Effective Altruism community, which has set itself the goal to determine where we can do the most good. Effective Altruism has adopted a common framework aligned with common scientific practices to determine how to prioritize between causes.”	Discourse by T	
24	elaboration & securing knowledge	With the help of the Socratic dialogue the T guides the S towards the five key aspects of Effective Altruism (scale, effectivity, neglectedness, alternative, expected value). T can for example ask: “If intervention one helps 1000 people and intervention two helps 10000, everything else equal, which intervention should be prioritized?” As soon as one of the key aspects of EA is identified by the S, it is written down by the S with a small description.	S-T-discussion	Writing material and or notebook

8.2.4 Lesson plan 4: Effective Altruism's biggest cause areas

time (min)	phase	content	arrangement	media
3	guidance	1-2 S repeat the five key aspects of EA discussed in the last lesson. T adds comments if necessary. T guides to the cause areas by saying: "After learning about which criteria help us determine which problems we should focus on, now let us continue to define those problems."	Discourse by S & eventual discourse by T	Notes by S
5	elaboration	The class is separated into five equally big groups. T writes down the biggest cause areas of EA (see Figure 5) in random order on the board. The groups are asked to order them by priority before disclosing the prioritization by the EA community. Correctly guessing of the position of one of the five cause areas gives the group one point. The more points a group obtains the earlier they are allowed to choose the cause area they shall elaborate on for a group work.	Discourse by T & group work	Figure 5 (Top causes in the EA community)
22	elaboration	The groups elaborate with the help of this thesis' chapter 6 (Effective Altruism's biggest cause areas) in the form of a handout and the internet a short presentation about the topic for the whole class.	Group work	Notebooks, web-enabled devices and handouts
15	securing knowledge	The groups present their cause areas to the class in descending order (see Figure 5). In case the elaboration of the topic took longer than expected, not all cause areas are presented.	Discourse by S	Depends on the presentation form each group chooses, eventually projector, notebook and or board and chalk

9 Conclusion

The aim of this thesis was to elaborate a case for presenting Effective Altruism in the ethics classes that are to be deployed in high school around Austria the school year of 2021/22. After careful evaluation, the conclusion of this thesis is that the topic of Effective Altruism perfectly suits the goals set out by the Federal Ministry for Education, Science and Research of Austria for ethics class. As a framework to deal with the complex life questions of altruism, it would certainly help students develop their reasoning and argumentation skills as well as learn new useful frameworks and mental tools to address real life challenges.

Effective Altruism is best understood as a global framework set to find the most pressing problems worldwide and determine how best to solve them. As MacAskill (2019, p. 14) put it, it applies scientific thinking (evidence-based decision making and careful reasoning) to make the most out of a given set of resources into the goal of maximizing their positive impact on others. This makes Effective Altruism not only a social movement, which has influenced thousands of people around the world and steered billions of dollars towards the most effective interventions, but also as a “[...] intellectual project (or ‘research field’) [...]” (MacAskill, 2019b, p. 14); an intellectual project in the search for the best solutions to the most pressing issues humanity is currently facing and might face in the future.

Doing so requires several tools and frameworks which the Effective Altruism community has adopted from the scientific field, e.g., expected value analyses, mistrust your intuitive judgements, as well as questioning one’s actions. Understanding that prioritization is essential to be able to help the biggest number of individuals, requires not only stepping back to look at the bigger picture but also to be able to change one’s focus towards a more pressing problem or readapting one’s approach. Many in the EA community do not necessarily work directly in the field, or at a research facility or think tank, but have decided to dedicate their lives to help those working there by providing them with funding. They have adopted high-earning careers to donate big shares of their income towards those causes, helping them do their work in the first place.

The questions Effective Altruism raises and the frameworks it fosters fit ideally with the goals the Federal Ministry for Education, Science and Research of Austria has set for the ethics classes that are to be established in high school as of the schoolyear 2021/22. The fostering of critical thinking, valid argumentation and reflection on ethical questions as well as the empowerment of students to deal with fundamental life questions are among the pre-defined

goals of the ethics class curriculum. Students are also to be encouraged to critically evaluate their own attitudes and moral standpoints and learn to find solutions to moral problems (BMBWF, 2020, p. 1). Effective Altruism seamlessly fits these goals as it deals with finding solutions for real life ethical global problems.

The lesson plans presented in this thesis provide a guidance for teachers to introduce the topic of Effective Altruism into ethics class. They are set out to guide the students through the subject, motivating them to engage with the complex questions Effective Altruism raises and critically evaluate their own moral judgements and positions in the process. At the end of the lesson series, students should be able to discuss global matters and reflect on the impact of their own actions on the situation of others around the world. They should also develop capabilities and skills that will enable them to better deal with complex issues they might face in their personal lives.

Effective Altruism is definitely a subject worth discussing with students in ethics class. It raises important life-changing questions, accentuating the impact every single one of us might have on the life of others. Being a citizen of the world requires us to confront the reality that the world we live in is far from perfect. There are problems, we might not be directly confronted with, but are nevertheless ethical issues we might should be addressing as if they were happening right in front of us. Effective Altruism opens up our eyes and gives us the tools to make the right questions and find the best solutions to solve the world's most pressing problems.

10 Summary

As a future teacher of ethics classes in Austria, I am concerned to discuss not only about the theoretical aspects of ethics with my future students but also about its practice. Effective Altruism is a growing social movement based on the idea to help others as much and as good as possible. The idea first developed from the ethical argument first presented by Peter Singer in 1972 that as individuals living middle class lives in the West, we have not only enough resources to help the global poor, but we have a certain responsibility to do so. It took some years until this idea and parallel efforts by several individuals looking for the best ways to help the global poor progressed to become the global movement it is today.

The focus of the movement has expanded to include other problems besides global poverty such as climate change, the risks of artificial intelligence, and non-human animal suffering. It is committed to find and assess the biggest and most urgent problems humanity faces and help develop solutions to fix them. Today, thousands of effective altruists – as the movements' members are called – not only donated and steered millions of dollars to evidence-based charities around the world but also to programs striving to develop promising solutions to current problems.

This thesis provides an overview of Effective Altruism, its history and key argumentation, its key concepts, its current areas of focus as well as its most highlighted critique points. It also discusses some of the most prominent criticism of the movement from academia and the charity sector. A short plan for discussing the topic in class is presented in each chapter and further elaborated in a chapter of its own (Chapter 8). After evaluation, the topic of Effective Altruism seems to seamlessly fit the ethics class curriculum set out by the Austrian Federal Ministry for Education, Science and Research. It raises questions that will introduce students into the most pressing problems worldwide, they might not necessarily be exposed to in their everyday life. Discussing EA with students will help them critically analyze their own moral attitudes towards global problems and teach them how their actions matter in the interconnected world we live in.

11 Abstract

EN

Effective Altruism is a growing global movement with the goal to find and implement solutions for the world's biggest problems. Its members address the question of how best to implement our resources to create the "most good possible" with a science-based approach. After elaboration and evaluation of the topic, Effective Altruism is deemed fitting with the preliminary curriculum for the high school ethics classes to be established in the schoolyear 2021/22. This thesis provides an overview of Effective Altruism and discusses one possible approach to introduce the topic into ethics classes in Austrian high schools.

DE

Der Effektive Altruismus ist eine wachsende globale Bewegung, die es sich zur Aufgabe gemacht hat, die effektivsten Lösungen für die größten Probleme der Menschheit zu finden. Mit Hilfe wissenschaftlicher Methoden und Denkansätze versuchen die Mitglieder herauszufinden, wie sie ihre Ressourcen am besten einsetzen sollten, um den größtmöglichen positiven Effekt zu erzielen. Nach intensiver Ausarbeitung und Evaluierung des Themas wurde der Effektive Altruismus als geeignetes Unterrichtsthema für den zukünftigen Ethikunterricht in der österreichischen Sekundarstufe bewertet. Der Effektive Altruismus scheint im Einklang mit dem vorläufigen Curriculum für den ab dem Schuljahr 2021/22 etablierten Ethikunterricht zu sein. Diese Masterarbeit bietet eine Einführung in das Thema des Effektiven Altruismus und einen möglichen Ansatz zur Einführung des Themas im Ethikunterricht in den Schulen Österreichs.

12 Bibliography

- 80,000 Hours. (n.d.). *Our current list of the most important world problems*. 80,000 Hours. Retrieved 18 March 2021, from <https://80000hours.org/problem-profiles>
- AGI Society. (n.d.). *AGI Society*. AGI Society. Retrieved 23 March 2021, from <http://www.agi-society.org>
- AidGrade. (n.d.). How is AidGrade Different? *AidGrade*. Retrieved 18 March 2021, from <http://www.aidgrade.org/how-is-aidgrade-different>
- Animal Charity Evaluators. (2016, November). Why Farmed Animals? *Animal Charity Evaluators*. <https://animalcharityevaluators.org/donation-advice/why-farmed-animals/>
- Ann Violin, Mary (1990). Pythagoras- The First Animal Rights Philosopher. *Between the Species*, 6(3). <https://doi.org/10.15368/bts.1990v6n3.6>
- Balter, Michael (2013, April 25). *Strongest Evidence of Animal Culture Seen in Monkeys and Whales*. Science | AAAS. <https://www.sciencemag.org/news/2013/04/strongest-evidence-animal-culture-seen-monkeys-and-whales>
- Banerjee, Abhijit & Duflo, Esther (2012). *Poor Economics: A Radical Rethinking of the Way to Fight Global Poverty* (Reprint edition). PublicAffairs.
- Banks, Antoine; Calvo, Ernesto; Karol, David & Telhami, Shibley (2020). #PolarizedFeeds: Three Experiments on Polarization, Framing, and Social Media. *The International Journal of Press/Politics*, 1940161220940964. <https://doi.org/10.1177/1940161220940964>
- Beard, Simon & Torres, Phil (2020). *Ripples on the Great Sea of Life: A Brief History of Existential Risk Studies*. SSRN. <http://dx.doi.org/10.2139/ssrn.3730000>
- Bedau, Hugo (1983). Bentham's Utilitarian Critique of the Death Penalty. *The Journal of Criminal Law and Criminology* (1973-), 74(3), 1033–1065. <https://doi.org/10.2307/1143143>
- Bentham, Jeremy (1789). *An Introduction to the Principles of Morals and Legislation* (J. Bennett, Ed.; New Edition published in 1823). <https://www.earlymoderntexts.com/assets/pdfs/bentham1780.pdf>
- Berger, Ken & Penna, Robert (2013, November 25). *The Elitist Philanthropy of So-Called Effective Altruism (SSIR)*. Stanford Social Innovation Review. https://ssir.org/articles/entry/the_elitist_philanthropy_of_so_called_effective_altruism
- Berkey, Brian (2018). The Institutional Critique of Effective Altruism. *Utilitas*, 30(2), 143–171. <https://doi.org/10.1017/S0953820817000176>
- BMBWF (2020). *Ethik – Pflichtgegenstand für alle Schülerinnen und Schüler, die keinen Religionsunterricht besuchen*. <https://www.bmbwf.gv.at/Themen/schule/zrp/ethik.html>
- Boeckel, Thomas; Brower, Charles; Gilbert, Marius; Grenfell, Bryan; Levin, Simon; Robinson, Timothy; Teillant, Aude & Laxminarayan, Ramanan (2015). Global trends in antimicrobial use in food animals. *Proceedings of the National Academy of Sciences*, 112(18), 5649–5654. <https://doi.org/10.1073/pnas.1503141112>
- Bostrom, Nick (2013). Existential Risk Prevention as Global Priority. *Global Policy*, 4(1), 15–31. <https://doi.org/10.1111/1758-5899.12002>

- Burson, Katherine; Smith, Robert & Faro, David (2011). Flocks, Herds, and Families: The Influence of Victim-Unitization on Charitable Giving. *ACR North American Advances*, NA-38. <https://www.acrwebsite.org/volumes/16137/volumes/v38/NA-38>
- CDC (2019a, August 2). *About MERS*. Centers for Disease Control and Prevention. <https://www.cdc.gov/coronavirus/mers/about/index.html>
- CDC (2019b, August 2). *MERS Frequently Asked Questions & Answers*. Centers for Disease Control and Prevention. <https://www.cdc.gov/coronavirus/mers/faq.html>
- CDC (2019c, November 20). *First Global Estimates of 2009 H1N1 Pandemic Mortality Released by CDC-Led Collaboration | CDC*. Centers for Disease Control and Prevention. <https://www.cdc.gov/flu/spotlights/pandemic-global-estimates.htm>
- CEA (n.d.). *CEA's Guiding Principles*. Centre for Effective Altruism. Retrieved 27 March 2021, from <https://www.centreforeffectivealtruism.org/ceas-guiding-principles/>
- CEA (2020, March 1). *Introduction to Effective Altruism*. Effective Altruism. <https://www.effectivealtruism.org/articles/introduction-to-effective-altruism/>
- Charity Navigator (n.d.). *Overview*. Charity Navigator. Retrieved 26 February 2021, from <http://www.charitynavigator.org/index.cfm?bay=content.view&cpid=628>
- Charity Navigator (2018). *Giving Statistics*. Charity Navigator. <http://www.charitynavigator.org/index.cfm?>
- Charity Navigator (2021). *Charity Navigator—Your Guide To Intelligent Giving | Home*. Charity Navigator. <http://www.charitynavigator.org/index.cfm?>
- Chittilapally, Joslyn (2020, April 15). *How factory farming breeds deadly viruses and epidemics*. LifeGate. <https://www.lifegate.com/factory-farming-epidemics-coronavirus>
- Collins, Chuck; Ocampo, Omar & Paslaski, Sophia (2020). *Billionaire Bonanza 2020*. Institute for Policy Studies. <https://ips-dc.org/wp-content/uploads/2020/04/Billionaire-Bonanza-2020.pdf>
- Cornejo-Juárez, Patricia; Vilar-Compte, Diana; Pérez-Jiménez, Carolina; Namendys-Silva, Silvio; Sandoval-Hernández, Silvia & Volkow-Fernández, Patricia (2015). The impact of hospital-acquired infections with multidrug-resistant bacteria in an oncology intensive care unit. *International Journal of Infectious Diseases: IJID: Official Publication of the International Society for Infectious Diseases*, 31, 31–34. <https://doi.org/10.1016/j.ijid.2014.12.022>
- Crimmins, James (2019). Jeremy Bentham. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (2019th ed.). Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/sum2020/entries/bentham/>
- Crompton, Louis (1978). Jeremy Bentham's essay on 'Paederasty'. Part 2. *Journal of Homosexuality*, 4(1), 91–107. https://doi.org/10.1300/J082v04n01_07
- de Castro, Marcia Caldas & Singer, Burton H. (2005). Was malaria present in the Amazon before the European conquest? Available evidence and future research agenda. *Journal of Archaeological Science*, 32(3), 337–340. <https://doi.org/10.1016/j.jas.2004.10.004>
- Desvousges, William H.; Johnson, Reed; Dunford, Richard W.; Hudson, Sara P.; Wilson, Nicole & Boyle, Kevin J. (1993). Chapter III - Measuring Natural Resource Damages with Contingent Valuation: Tests of Validity and Reliability. In J. A. Hausman (Ed.),

- Contributions to Economic Analysis* (Vol. 220, pp. 91–164). Elsevier. <https://doi.org/10.1016/B978-0-444-81469-2.50009-2>
- Diogenes (2016, January 20). *EA is elitist. Should it stay that way?* - *EA Forum*. <https://forum.effectivealtruism.org/posts/YLaw4xHb39kZCdsAS/ea-is-elitist-should-it-stay-that-way>
- Duda, Roman (2018, July). *How would you spend \$500,000,000,000?* 80,000 Hours. <https://80000hours.org/problem-profiles/global-priorities-research/>
- Dullaghan, Neil (2019, December 5). *EA Survey 2019 Series: Community Demographics & Characteristics* - *EA Forum*. Effective Altruism Forum. <https://forum.effectivealtruism.org/posts/wtQ3XCL35uxjXpwjE/ea-survey-2019-series-community-demographics-and>
- Dunayer, Joan (2004). *Speciesism*. Ryce Publishing
- EA Forum (n.d.). *Criticism (EA Cause Areas)*—*EA Forum*. Effective Altruism Forum. Retrieved 17 March 2021, from <https://forum.effectivealtruism.org/tag/criticism-ea-cause-areas>
- Easterly, William (2007). *The White Man's Burden: Why the West's Efforts to Aid the Rest Have Done So Much Ill and So Little Good* (Annotated edition). Penguin Books.
- EffectiveAltruism.org (n.d.). *Counterfactual considerations*. Effective Altruism Concepts. Retrieved 2 March 2021, from <https://concepts.effectivealtruism.com/concepts/counterfactual-considerations/>
- EffectiveAltruism.org (2020, March 1). *Introduction to Effective Altruism*. Effective Altruism. <https://www.effectivealtruism.org/articles/introduction-to-effective-altruism/>
- Evidence Action (2020, August 4). *The 44 Cent Treatment That Dramatically Changed Lives in Kenya*. *Evidence Action*. <https://www.evidenceaction.org/press-release-the-44-cent-treatment-that-dramatically-changed-lives-in-kenya/>
- Evidence Action (2021). *Deworm the World*. *Evidence Action*. <https://www.evidenceaction.org/dewormtheworld/>
- FHI (n.d.). *Future of Humanity Institute*. The Future of Humanity Institute. Retrieved 23 March 2021, from <http://www.fhi.ox.ac.uk/>
- fishcount.org (2019a). *Fishcount estimates of numbers of individuals killed in (FAO) reported fishery production*. http://fishcount.org.uk/studydatascreens/2016/fishcount_estimates_list.php
- fishcount.org (2019b). *Numbers of fish caught from the wild each year* | *fishcount.org.uk*. <http://fishcount.org.uk/fish-count-estimates-2/numbers-of-fish-caught-from-the-wild-each-year>
- Food and Drug Administration (2014). *Summary Report On Antimicrobials Sold or Distributed for Use in Food-Producing Animals*. Food and Drug Administration. <http://www.fda.gov/downloads/ForIndustry/UserFees/AnimalDrugUserFeeActADUFA/UCM338170.pdf>.
- Friedman, Eric (2013). *Reinventing Philanthropy: A Framework for More Effective Giving*. POTOMAC BOOKS.
- Gabriel, Iason (2017). Effective Altruism and its Critics. *Journal of Applied Philosophy*, 34(4), 457–473. <https://doi.org/10.1111/japp.12176>

- Gabrieli, Nick (2021, February). *Nick Gabrieli: Effective Altruism and Theories of Welfare: a Philosophical Analysis*. Effective Altruism at Harvard. <https://www.harvardea.org/blog/2021/2/15/nick-gabrieli-effective-altruism-and-theories-of-welfare-a-philosophical-analysis>
- Gamborg, Christian; Palmer, Clare & Sandoe, Peter (2012). Ethics of Wildlife Management and Conservation: What Should We Try to Protect? | Learn Science at Scitable. *Nature Education Knowledge*, 3(10), 8. <https://www.nature.com/scitable/knowledge/library/ethics-of-wildlife-management-and-conservation-what-80060473/>
- GCRI (n.d.). *About | Global Catastrophic Risk Institute*. Global Catastrophic Risk Institute. Retrieved 23 March 2021, from <https://gcrinstitute.org/about/>
- Gertler, Aaron (2019, November 12). *EA Leaders Forum: Survey on EA priorities (data and analysis)—EA Forum*. Effective Altruism Forum. https://forum.effectivealtruism.org/posts/TpoeJ9A2G5Sipxfit/ea-leaders-forum-survey-on-ea-priorities-data-and-analysis%23Cause_Priorities#Known_causes
- GiveWell (n.d.-a). *GiveWell | Charity Reviews and Research*. GiveWell. Retrieved 18 March 2021, from <https://www.givewell.org/>
- GiveWell (n.d.-b). *Other Charity Reviews*. GiveWell. Retrieved 18 March 2021, from <https://www.givewell.org/research/other-charity-reviews>
- GiveWell (n.d.-c). *Your Dollar Goes Further Overseas*. GiveWell. Retrieved 26 February 2021, from <https://www.givewell.org/giving101/Your-dollar-goes-further-overseas>
- GiveWell (2016, November). *Process for Identifying Top Charities*. GiveWell. <https://www.givewell.org/how-we-work/process>
- GiveWell (2018, February). *Our Story*. GiveWell. <https://www.givewell.org/about/story>
- GiveWell (2020a, October). *Our Mistakes*. GiveWell. <https://www.givewell.org/about/our-mistakes>
- GiveWell (2020b, November). *Evidence Action's Deworm the World Initiative*. GiveWell. <https://www.givewell.org/charities/deworm-world-initiative>
- GiveWell (2020c, November). *Malaria Consortium – Seasonal Malaria Chemoprevention*. GiveWell. <https://www.givewell.org/charities/malaria-consortium>
- GiveWell (2020d, November). *Our Top Charities | GiveWell*. GiveWell. <https://www.givewell.org/charities/top-charities>
- GiveWell (2020e, December). *GiveWell's Impact*. GiveWell. <https://www.givewell.org/about/impact>
- Global Preparedness Monitoring Board (2019). *A World at Risk. Annual report on global preparedness for health emergencies*. (p. 48) [Annual Report]. Global Preparedness Monitoring Board. https://apps.who.int/gpmb/assets/annual_report/GPMB_annualreport_2019.pdf
- Good, Irving J. (1966). Speculations Concerning the First Ultraintelligent Machine. *Advances in Computers*, 6, 31–88. [https://doi.org/10.1016/S0065-2458\(08\)60418-0](https://doi.org/10.1016/S0065-2458(08)60418-0)
- Good Ventures (n.d.). *About Us*. Good Ventures. Retrieved 18 March 2021, from <https://www.goodventures.org/about-us>

- GWWC (2021). *What are the best charities to donate to in 2021?* Giving What We Can. <https://www.givingwhatwecan.org/best-charities-to-donate-to-2021/>
- Hamory, Joan; Miguel, Edward; Walker, Michael; Kremer, Michael & Baird, Sarah J. (2020). *Twenty Year Economic Impacts of Deworming*. <https://doi.org/10.3386/w27611>
- Hassenfeld, Elie (2007, November 28). *Cheap ways to save lives*. The GiveWell Blog. <https://blog.givewell.org/2007/11/27/cheap-ways-to-save-lives/>
- Heard, Andrew (1997). *The Challenge of Utilitarianism and Relativism to Human Rights: Chimeras in Sheep's Clothing?* <http://www.sfu.ca/~aheard/417/util.html>
- Henderson, Donald A. (2011). The eradication of smallpox – An overview of the past, present, and future. *Vaccine*, 29, D7–D9. <https://doi.org/10.1016/j.vaccine.2011.06.080>
- Henig, Robin M. (2020, August 4). *Experts warned of a pandemic decades ago. Why weren't we ready?* National Geographic. <https://www.nationalgeographic.com/science/article/experts-warned-pandemic-decades-ago-why-not-ready-for-coronavirus>
- Herzog, Lisa (2016, February 22). *Can 'effective altruism' really change the world?* OpenDemocracy. <https://www.opendemocracy.net/en/transformation/can-effective-altruism-really-change-world/>
- Hyde, Susan D. (2020). Democracy's backsliding in the international environment. *Science*, 369(6508), 1192–1196. <https://doi.org/10.1126/science.abb2434>
- Isaak, Jim & Hanna, Mina J. (2018). User Data Privacy: Facebook, Cambridge Analytica, and Privacy Protection. *Computer*, 51(8), 56–59. <https://doi.org/10.1109/MC.2018.3191268>
- Joey (2019, September 1). *Cause X Guide—EA Forum*. Effective Altruism Forum. <https://forum.effectivealtruism.org/posts/kFmFLcdSFKo2GFJkc/cause-x-guide>
- Karnofsky, Holden (2016, July 25). *Effective altruism*. The GiveWell Blog. <https://blog.givewell.org/2013/08/13/effective-altruism/>
- Kharas, Homi (2017). *The Unprecedented Expansion of the Global Middle Class—An Update* (No. 100). Global Economy and Development at Brookings.
- Kniess, Johannes (2019). Bentham on animal welfare. *British Journal for the History of Philosophy*, 27(3), 556–572. <https://doi.org/10.1080/09608788.2018.1524746>
- Lichtenberg, Judith (2015, November 30). Peter Singer's Extremely Altruistic Heirs. *The New Republic*. <https://newrepublic.com/article/124690/peter-singers-extremely-altruistic-heirs>
- Low, Philip (2012). *The Cambridge Declaration on Consciousness* (J. Panksepp, D. Reiss, D. Edelman, B. Van Swinderen, P. Low, & C. Koch, Eds.). Cambridge University Press. <https://fcmconference.org/img/CambridgeDeclarationOnConsciousness.pdf>
- MacAskill, William (2016). *Doing Good Better: How Effective Altruism Can Help You Help Others, Do Work that Matters, and Make Smarter Choices about Giving Back* (Reprint edition). Avery.
- MacAskill, William (2019a, July 19). *Defining Effective Altruism—EA Forum* [EA Forum]. <https://forum.effectivealtruism.org/posts/9wYa8BqSTMcx9j2tK/defining-effective-altruism>

- MacAskill, William (2019b). The Definition of Effective Altruism. In H. Greaves & T. Pummer, *Effective Altruism: Philosophical Issues*. Oxford University Press.
- Mackie, John D. (1977). *Ethics: Inventing Right and Wrong* (Reprint edition). Penguin Books.
- Make A Wish (2018). *Frequently Asked Questions*. Make A Wish. <https://wish.org/faq>
- Matta, Roula; Hallit, Souheil; Hallit, Rabih; Bawab, Wafaa; Rogues, Anne-Marie & Salameh, Pascale (2018). Epidemiology and microbiological profile comparison between community and hospital acquired infections: A multicenter retrospective study in Lebanon. *Journal of Infection and Public Health*, 11(3), 405–411. <https://doi.org/10.1016/j.jiph.2017.09.005>
- Miguel, Edward & Kremer, Michael (2003). Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities. *Econometrica*, 72(1), 159–217. <https://doi.org/10.1111/j.1468-0262.2004.00481.x>
- MIRI (n.d.). *Research*. Machine Intelligence Research Institute. Retrieved 22 March 2021, from <https://intelligence.org/research/>
- Monroe, Kristen R. (2001). Altruism and Self-interest. In N. J. Smelser & P. B. Baltes (Eds.), *International Encyclopedia of the Social & Behavioral Sciences* (pp. 415–418). Pergamon. <https://doi.org/10.1016/B0-08-043076-7/01093-7>
- Moses, Strydom & Sheryl, Buckley (2019). *AI and Big Data's Potential for Disruptive Innovation*. IGI Global.
- Moyo, Dambisa (2010). *Dead Aid: Why Aid Is Not Working and How There Is a Better Way for Africa* (Reprint edition). Farrar, Straus and Giroux.
- National Archives (2016, August 15). *Benjamin Franklin's Anti-Slavery Petitions to Congress*. National Archives. <https://www.archives.gov/legislative/features/franklin>
- National Intelligence Council (2020). *Near Future: Tensions are Rising*. Global Trends. <https://www.dni.gov/index.php/global-trends/near-future>
- Nguyen, Trang (2008). Information, role models and perceived returns to education experimental evidence from Madagascar. *Massachusetts Institute of Technology*. <https://www.povertyactionlab.org/sites/default/files/documents/Nguyen%202008.pdf>
- Noelle-Neumann, Elisabeth (2001). *Die Schweigespirale. Öffentliche Meinung—Unsere soziale Haut*. Langen/Müller.
- OECD (2020). *Income Distribution Database*. <https://stats.oecd.org/Index.aspx?DataSetCode=IDD>
- oesterreich.gv.at (2020, January 1). *Religionsunterricht*. oesterreich.gv.at - Österreichs digitales Amt. https://www.oesterreich.gv.at/themen/leben_in_oesterreich/kirchenein__austritt_und__religionen/Seite.820014.html
- OIF (2019). *Islamischer Religionsunterricht in Österreich. Fact Sheet* (Fact-Sheet No. 32). Österreichischer Integrationsfonds. https://www.integrationsfonds.at/fileadmin/user_upload/FactSheet-32_IslamRelunterr_1-6_SCREEN.pdf
- Open Philanthropy (2015, March 26). *Cause Selection*. Open Philanthropy. <https://www.openphilanthropy.org/research/cause-selection>

- Open Philanthropy (2016, July 19). *What Do We Know about AI Timelines?* Open Philanthropy. <https://www.openphilanthropy.org/focus/global-catastrophic-risks/potential-risks-advanced-artificial-intelligence/ai-timelines>
- Open Philanthropy (2021). *Progress to Date*. Open Philanthropy. <https://www.openphilanthropy.org/about/progress-to-date>
- O'Reilly, Michelle; Dogra, Nisha; Whiteman, Natasha; Hughes, Jason; Eruyar, Seyda & Reilly, Paul (2018). Is social media bad for mental health and wellbeing? Exploring the perspectives of adolescents. *Clinical Child Psychology and Psychiatry*, 23(4), 601–613. <https://doi.org/10.1177/1359104518775154>
- Our World in Data (2018). *Number of animals slaughtered for meat*. Our World in Data. <https://ourworldindata.org/grapher/animals-slaughtered-for-meat>
- Our World in Data (2021a). *Weekly new ICU admissions for COVID-19 per million*. Our World in Data. <https://ourworldindata.org/grapher/weekly-icu-admissions-covid-per-million>
- Our World in Data (2021b, March 11). *Daily vs. Cumulative confirmed deaths due to COVID-19*. Our World in Data. <https://ourworldindata.org/grapher/confirmed-covid-19-deaths-cumulative-vs-daily>
- Piper, Kelsey (2018, December 21). *The case for taking AI seriously as a threat to humanity*. Vox. <https://www.vox.com/future-perfect/2018/12/21/18126576/ai-artificial-intelligence-machine-learning-safety-alignment>
- Piper, Kelsey (2020, August 13). *GPT-3, explained: This new language AI is uncanny, funny — and a big deal*. Vox. <https://www.vox.com/future-perfect/21355768/gpt-3-ai-openai-turing-test-language>
- Piperaki, Evangelia T. & Daikos, Georgios L. (2016). Malaria in Europe: Emerging threat or minor nuisance? *Clinical Microbiology and Infection*, 22(6), 487–493. <https://doi.org/10.1016/j.cmi.2016.04.023>
- Riedel, Stefan (2005). Edward Jenner and the history of smallpox and vaccination. *Proceedings (Baylor University Medical Center)*, 18(1), 21–25. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1200696/>
- Ritchie, Hannah & Roser, Max (2019). *Causes of Death*. Our World in Data. <https://ourworldindata.org/causes-of-death>
- Roser, Max (2013). *Global Economic Inequality*. Our World in Data. <https://ourworldindata.org/global-economic-inequality>
- Roser, Max; Ortiz-Ospina, Esteban & Ritchie, Hannah (2019). *Life Expectancy*. Our World in Data. <https://ourworldindata.org/life-expectancy>
- Roser, Max & Ritchie, Hannah (2013). *Maternal Mortality*. Our World in Data. <https://ourworldindata.org/maternal-mortality>
- Roser, Max & Ritchie, Hannah (2016). *Burden of Disease*. Our World in Data. <https://ourworldindata.org/burden-of-disease>
- Roser, Max; Ritchie, Hannah & Dadonaite, Bernadeta (2019). *Child and Infant Mortality*. Our World in Data. <https://ourworldindata.org/child-mortality>
- Rudolph, Jared (2008). Consequences and Limits: A Critique of Consequentialism. *Macalester Journal of Philosophy*, 17(1), 64–76. <http://digitalcommons.macalester.edu/philo/vol17/iss1/12>

- Sassi, Franco (2006). Calculating QALYs, comparing QALY and DALY calculations. *Health Policy and Planning*, 21(5), 402–408. <https://doi.org/10.1093/heapol/czl018>
- Seixas, Brayan V. (2017). Welfarism and extra-welfarism: A critical overview. *Cadernos de Saúde Pública*, 33(8). <https://doi.org/10.1590/0102-311x00014317>
- Singer, Peter (2009). *Animal Liberation: The Definitive Classic of the Animal Movement* (Updated ed. edition). Harper Perennial Modern Classics.
- Singer, Peter (2010). *The Life You Can Save: How to Do Your Part to End World Poverty* (Reprint edition). Random House Trade Paperbacks.
- Singer, Peter (2015). *The Most Good You Can Do: How Effective Altruism Is Changing Ideas About Living Ethically* (1st edition). Yale University Press.
- Sinnott-Armstrong, Walter (2019). *Consequentialism*. The Stanford Encyclopedia of Philosophy. <https://plato.stanford.edu/archives/sum2019/entries/consequentialism/>
- Sung, Hyuna; Ferlay, Jacques; Siegel, Rebecca L.; Laversanne, Mathieu; Soerjomataram, Isabelle; Jemal, Ahmedin & Bray, Freddie (2021). Global cancer statistics 2020: GLOBOCAN estimates of incidence and mortality worldwide for 36 cancers in 185 countries. *CA: A Cancer Journal for Clinicians*, n/a(n/a). <https://doi.org/10.3322/caac.21660>
- Taleb, Nassim N. (2010). *The Black Swan: Second Edition: The Impact of the Highly Improbable: With a new section: 'On Robustness and Fragility'* (2nd ed. edition). Random House Trade Paperbacks.
- The World Bank (n.d.-a). *Poverty* [Text/HTML]. World Bank. Retrieved 8 February 2021, from <https://www.worldbank.org/en/topic/poverty>
- The World Bank (n.d.-b). *Poverty | Data*. Poverty. Retrieved 8 February 2021, from <https://data.worldbank.org/topic/poverty>
- The World Bank (n.d.-c). *Poverty—Overview* [Text/HTML]. World Bank. Retrieved 8 February 2021, from <https://www.worldbank.org/en/topic/poverty/overview>
- The World Bank (2012). *What Matters Most for School Health and School Feeding: A Framework Paper* (Saber No. 3; Working Paper Series). The World Bank. http://wbfiles.worldbank.org/documents/hdn/ed/saber/supporting_doc/Background/S_HN/Framework_SABER-School_Health.pdf
- TLYCS (2021). *The Life You Can Save—Best Charities for Effective Giving*. The Life You Can Save. <https://www.thelifeyoucansave.org/>
- Todd, Benjamin; Farquhar, Sebastian & Mills, Peter (2012, May). The Ethical Careers Debate: A Discussion between Ben Todd, Sebastian Farquhar, and Pete Mills, edited by Tom Cutterham. *The Oxford Left Review*, 7. <https://oxfordleftreview.files.wordpress.com/2012/07/issue-7-new.pdf>
- Todd, Benjamin (2014, January 21). *Which cause is most effective?* 80,000 Hours. <https://80000hours.org/2014/01/which-cause-is-most-effective-300/>
- Todd, Benjamin (2017, October). *The case for reducing existential risks*. 80,000 Hours. <https://80000hours.org/articles/existential-risks/>
- Tomasik, Brian (2009). *How Many Wild Animals Are There? Last update on 07.08.2019*. Reducing Suffering. <https://reducing-suffering.org/how-many-wild-animals-are-there/>

- Tomasik, Brian (2015). The Importance of Wild-Animal Suffering. *Relations. Beyond Anthropocentrism*, 3(2), 133–152. <https://doi.org/10.7358/rela-2015-002-toma>
- Tomasik, Brian (2017, August 14). *Suffering in Animals vs. Humans*. Reducing Suffering. <https://reducing-suffering.org/suffering-in-animals-vs-humans/>
- Vinding, Magnus (2015). *Speciesism: Why It Is Wrong and the Implications of Rejecting It*.
- Vinding, Magnus (2018). *Effective Altruism: How Can We Best Help Others?* CreateSpace Independent Publishing Platform.
- Walker, Brandon (2020, March 15). *The Games That AI Won*. Medium. <https://towardsdatascience.com/the-games-that-ai-won-ff8fd4a71efc>
- Weissman, Gideon (2015, April 21). *Even if you don't eat meat, antibiotic resistance in food animals is a threat to your health*. Frontier Group. <https://frontiergroup.org/blogs/blog/fg/even-if-you-don%E2%80%99t-eat-meat-antibiotic-resistance-food-animals-threat-your-health>
- WHO (2016, February 10). *WHO | World's largest deworming campaign targets 270 million children in one day*. WHO; World Health Organization. http://www.who.int/neglected_diseases/news/270_million_children_treated_in_one_day/en/
- WHO (2019a, April). *WHO | Short overview of the malaria vaccine implementation programme*. WHO; World Health Organization. <http://www.who.int/malaria/media/malaria-vaccine-overview/en/>
- WHO (2019b, July 8). *WHO | Worldwide reduction in MERS cases and deaths since 2016*. World Health Organization; World Health Organization. <http://www.who.int/emergencies/mers-cov/worldwide-reduction-in-mers-cases-and-deaths-since-2016/en/>
- WHO (2020, November 30). *Fact sheet about Malaria*. World Health Organization. <https://www.who.int/news-room/fact-sheets/detail/malaria>
- WHO (2021, February 26). *WHO | Human infection with avian influenza A (H5N8) – the Russian Federation*. World Health Organization; World Health Organization. <http://www.who.int/csr/don/26-feb-2021-influenza-a-russian-federation/en/>
- Wiblin, Robert (2015, July 16). *Disagreeing about what's effective isn't disagreeing with effective altruism*. 80,000 Hours. <https://80000hours.org/2015/07/disagreeing-about-whats-effective-isnt-disagreeing-with-effective-altruism/>
- Wiblin, Robert (2017a, March). *Positively shaping the development of artificial intelligence*. 80,000 Hours. <https://80000hours.org/problem-profiles/positively-shaping-artificial-intelligence/>
- Wiblin, Robert (2017b, April 6). *How accurately does anyone know the global distribution of income?* 80,000 Hours. <https://80000hours.org/2017/04/how-accurately-does-anyone-know-the-global-distribution-of-income/>
- Wiblin, Robert; Koehler, Arden & Harris, Keiran (2019, December 5). *Peter Singer on being provocative, how his moral views have changed, and saving children drowning in shallow ponds*. 80,000 Hours. <https://80000hours.org/podcast/episodes/peter-singer-advocacy-and-the-life-you-can-save/>
- Wild Animal Initiative (n.d.). *About us*. Wild Animal Initiative. Retrieved 21 March 2021, from <https://www.wildanimalinitiative.org/about-us>

- Wild-Animal Suffering Research (n.d.). *Wild-Animal Suffering Research*. Wild-Animal Suffering Research. Retrieved 21 March 2021, from <https://was-research.org/>
- Williford, Miriam (1975). Bentham on the Rights of Women. *Journal of the History of Ideas*, 36(1), 167–176. <https://doi.org/10.2307/2709019>
- Wooster, Martin M. (2017, September 8). *My problem with Effective Altruism*. Philanthropy Daily. <https://www.philanthropydaily.com/my-problem-with-effective-altruism/>
- World Health Organization (2020, September 8). *Children: Improving survival and well-being*. World Health Organization. <https://www.who.int/news-room/fact-sheets/detail/children-reducing-mortality>
- World Population Review (2021). *Richest Countries In The World 2021*. World Population Review. <https://worldpopulationreview.com/country-rankings/richest-countries-in-the-world>
- World Vision (2019, November 25). 2010 Haiti earthquake: Facts, FAQs, and how to help. *World Vision*. <https://www.worldvision.org/disaster-relief-news-stories/2010-haiti-earthquake-facts>
- WWF (n.d.). *Australian bushfires: How your donation is making a difference*. WWF. Retrieved 21 March 2021, from <https://www.wwf.org.au/what-we-do/bushfire-recovery/how-your-donation-is-making-a-difference/how-your-donation-is-making-a-difference>
- You, Yaqi & Silbergeld, Ellen K. (2014). Learning from agriculture: Understanding low-dose antimicrobials as drivers of resistome expansion. *Frontiers in Microbiology*, 5. <https://doi.org/10.3389/fmicb.2014.00284>