The **Effective** Altruism Handbook



Edited by **Ryan Carey**

The Effective Altruism Handbook

The Effective Altruism Handbook

Edited by Ryan Carey



Published in 2015 by the Centre For Effective Altruism Oxford, Oxfordshire United Kingdon centreforeffectivealtruism.org

New and reprinted essays in this compilation are the property of the original author or publisher and are published here with their kind permission.

Thanks also to Peter Orr, who helped substantially by copyediting this book.

Contents

Intr	roduction	viii
I	WHAT IS EFFECTIVE ALTRUISM?	
1	The Drowning Child and the Expanding Circle	
	Peter Singer	3
2	What is Effective Altruism?	
	William MacAskill	11
3	Scope Neglect	
	Eliezer Yudkowsky	14
4	Tradeoffs	
	Julia Wise	19
II	CHARITY EVALUATION	
5	Efficient Charity: Do Unto Others	
	Scott Alexander	25
6	"Efficiency" Measures Miss the Point	
	Dan Pallotta	32
7	How Not to Be a "White in Shining Armor"	
	Holden Karnofsky	35
8	Estimation Is the Best We Have	
	Katja Grace	39
9	Our Updated Top Charities	
	Elie Hassenfeld	42
III	CAREER CHOICE	
10	Don't Get a Job at a Charity: Work on Wall Street	
	William MacAskill	67

Contents

11	High Impact Science	
	Carl Shulman	71
12	How to Assess the Impact of a Career	
	Ben Todd	76
IV	CAUSE SELECTION	
13	Your Dollar Goes Further Overseas	
	GiveWell	83
14	The Haste Consideration	
	Matt Wage	85
15	Preventing Human Extinction	
	Nick Beckstead, Peter Singer & Matt Wage	88
16	Speciesism	
	Peter Singer	95
17	Four Focus Areas of Effective Altruism	
	Luke Muehlhauser	99
V	ORGANIZATIONS	
18	GiveWell	
	GiveWell	109
19	Giving What We Can	
	Michelle Hutchinson, Executive Director	114
20	The Life You Can Save	
	Charlie Bresler, Executive Director	117
21	80,000 Hours	
	Ben Todd, Co-founder and Executive Director	120
22	Charity Science	
	Xiomara Kikauka, Co-founder	124
23	Effective Altruism Foundation	
	Adriano Mannino, President and Jonas Vollmer, Execu-	
	tive Director	127
24	The Machine Intelligence Research Institute	

	Luke Muehlhauser, Executive Director	131
25	Animal Charity Evaluators	
	Jon Bockman, Executive Director	134
Afte	erword	139

Introduction

Peter Singer and William MacAskill

Welcome to the effective altruism community! You're about to find out about what we think is the world's most exciting and important new social movement.

What is effective altruism?

Effective altruism is a growing community based around the idea of aiming to do the most good that one can. It involves:

- 1. Being open to all the possible ways to do good and pursuing the path with the biggest positive impact;
- 2. Using evidence to figure out how to do the most good; and
- 3. Choosing to make altruism a significant part of one's life.

What does that mean?

You have approximately 80,000 working hours in your life. While this might feel like a lot, it's not much when compared to the scale of the problems in the world, such as extreme

poverty, global ill-heath, climate change, gender inequality, animal suffering, risk of war, and thousands more.

These are all pressing problems, and there's simply no way that one person can solve *all* of these problems in their lifetime. This means we need to prioritize, and try to answer the question: of all the possible ways of using one's time on earth to make a difference, which make the *most* difference? That's the central question that effective altruists ask themselves.

The question is extremely difficult, but there is a growing body of research that can help us to answer it:

- GiveWell searches for the charities that do the most good with every pound or dollar they receive.
- The Poverty Action Lab and Innovations for Poverty Action undertake detailed scientific investigations in order to assess different development interventions.
- Animal Charity Evaluators research charities that seek to reduce animal suffering and recommends those that it believes are most cost-effective.
- The Copenhagen Consensus Centre gathers expert opinion from economists in order to assess which activities have the best cost-benefit ratio.
- The Global Priorities Project does foundational research on comparing the cost-effectiveness of different causes.

 80,000 Hours provides advice and research to help people choose a career where they can maximize their impact.

Through this research, we have found that it's in the power of individuals to do a tremendous amount of good. Simply by pledging to donate a proportion of one's earnings to the most cost-effective causes, or by altering one's career trajectory, we each have the power to significantly improve the lives of thousands of others.

Collectively, the effective altruism community has accomplished incredible things in just a few short years. Through Giving What We Can, over \$250 million has been pledged to the most cost-effective charities. More than 16,000 people have pledged through The Life You Can Save to give a portion of their income to effective charities helping the global poor. GiveWell has already raised over \$30 million, and now advises GoodVentures, a \$3 billion foundation. That means that we've bought and distributed over one and a half million bed nets, treated over six million children for parasitic worm infections, and provided over \$10 million in direct cash transfers to the world's poorest people, and we expect these numbers to grow significantly in the future.

As well as money raised, we've seen impact through a variety of other means. Hundreds of young people are choosing careers in order to maximize their impact, often based on the advice that 80,000 Hours provides. Researchers at the

Centre for Effective Altruism are involved in providing policy advice to the UK government. Effective Altruist ideas have been featured in every major US and UK news source, and through other avenues around the world. And a number of new non-profit and for-profit social enterprises have been set up, already with impressive results.

What do effective altruists do?

There are a wide variety of ways to do a huge amount of good in the world.

Many effective altruists do good by making donations. For example, Giving What We Can members like Aveek Bhattarchya and Catriona Mackay give 10% of their annual earnings to the most effective charities. Julia Wise, a social worker and her husband Jeff Kaufman, a software developer, donate 50% of their income to the most cost-effective causes.

Here are some examples of other paths:

- Setting up companies with a social mission. Lincoln Quirk co-founded Wave, which dramatically reduces the costs for US immigrants to send remittances back to their families.
- Setting up non-profits. Finding that university wasn't teaching them the skills they needed, Joey Savoie and Xio Kikauka left in order to set up Charity Science, which fundraises for the most cost-effective charities.

- Research. Holden Karnofsky and Elie Hassenfeld quit their jobs at the hedge fund Bridgewater in order to found GiveWell and evaluate charities.
- Politics. Habiba Islam took up law and is now building skills in strategy consulting with the aim of moving into politics.
- Advocacy and education. Julia Galef and Anna Salamon founded the Centre for Applied Rationality, in order to improve the decision-making skills of businesspeople, policymakers and non-profit organizations.

Where did the idea of effective altruism come from?

The idea of effective altruism arose naturally out of recent developments in economics, psychology and moral philosophy, in particular:

- The rise in evidence-based development aid, such as the use of randomized controlled trials led by economists at the Poverty Action Lab. These provide high-quality research about what works and what doesn't in development aid.
- The development of the heuristics and biases literature by psychologists Daniel Kahneman and Amos Tversky.
 This literature shows the failures of human rationality,

and thereby opens up the possibility of increasing one's impact by deliberately countering these biases.

 The development of moral arguments, by Peter Singer and others, in favor of there being a duty to use a proportion of one's resources to fight global poverty, and in favor of an "expanded moral circle" that gives moral weight to distant strangers, future people and nonhuman animals.

The effective altruism community grew out of a few different groups. GiveWell had been investigating charity effectiveness since 2007, and those in the rationalist community LessWrong had discussed how one could correct human biases in order to do more good since 2006.

In 2009, Giving What We Can was launched in Oxford by Toby Ord and William MacAskill, with the aim of forming a community based around devoting a significant proportion of one's resources to the causes that do the most good. Because of Toby's pledge to give everything he earned above \$28,000 per year, Giving What We Can generated significant media attention. This attracted like-minded people from across the globe: people who had had similar ideas but were comparatively isolated from each other. The community then grew rapidly as a result of interest from GiveWell donors, LessWrong, and followers of Peter Singer's work.

At the same time, people became increasingly interested in optimizing not just their charitable donations, but their lives in general. This led to the founding of 80,000 Hours in February 2011, which solidified the idea of a *movement* based around these ideas.

The label "effective altruism" became official when the Centre for Effective Altruism was founded in December 2011. The use of the term increased considerably when Peter Singer's TED Talk on the topic went live in May 2013.

How can I get involved?

The easiest way to start having an impact is by setting up a regular donation to one of GiveWell's top-recommended charities. Even if you don't have a lot to give, it's a good habit to get into, and means that you will start improving people's lives right away.

The easiest way to get more involved in the community is to introduce yourself on the Facebook Group, and ask to Skype someone. There's a lot going on, and it's easiest to get to know things by speaking to someone directly. Because we all have the same aims, we're very willing to help each other out, and we always enjoy talking about ways of increasing our impact.

Beyond that, you might want to consider taking Giving What We Can's pledge to give 10%, or The Life You Can Save's pledge (which varies according to your income), or you could sign up for a one-on-one career coaching session at 80,000

Hours, or start reading more about effective altruism — see below!

The rest of this collection

This handbook is a collection of blog posts and short articles that have been written by people who are engaged in effective altruism.

This collection has four parts. The first part introduces in more detail some of the ideas behind effective altruism. Peter Singer introduces the key argument for the idea that we are morally obliged to devote a significant proportion of our resources to helping others. William MacAskill discusses how effective altruism differs from conventional ways of thinking about helping others. Eliezer Yudkowsky describes one way in which the human brain is biased, the result of which is that we should expect ourselves to be poor at making good altruistic decisions. Julia Wise writes about trading altruistic and personal expenditure. She says that being an effective altruist is not a matter of self-sacrifice and grudging drudgery. Rather, it's about being excited about making the world a better place, and inspiring others to do the same.

The next three parts focus on the three most important questions that we face: (i) Charity Evaluation: How should we spend our *money*?; (ii) Career Choice: How should we spend our *time*?; and (iii) Cause Selection: Which *cause* should we pursue?

On Charity Evaluation, Scott Alexander introduces the idea of "efficient charity" — trying to maximize the impact of one's donations. Dan Pallotta criticizes the most commonly used metric for assessing charities — the "overheads ratio". Holden Karnofsky raises the general issue of trying to deal with problems of which one has little initial knowledge. Katja Grace discusses the difficulty and importance of trying to quantify the benefits that one produces. The chapter ends with Elie from GiveWell recommending their current top charities.

On Career Choice, William MacAskill introduces the key arguments behind the idea of *earning to give*. Carl Shulman describes the potential for scientific research to do a huge amount of good. The chapter ends with Ben Todd describes a framework for assessing the social impact of different careers and lists 80,000 Hours' current top-recommended career paths.

In Cause Selection, GiveWell provides reasons to alleviate poverty in developing countries rather than rich countries. Peter Singer argues that animal welfare should be considered a priority. Matt Wage gives a consideration in favor of growing the effective altruist movement. Nick Beckstead, Peter Singer, and Matt Wage give the argument that mitigating existential risk might be the most important cause. Luke Muehlhauser summarizes the four kinds of causes that effective altruists tend to focus on.

In the last section, each of the effective altruist organizations draws on the ideas from the preceding chapters to suggest what concrete actions you can take to make the world better.

We hope you enjoy learning more!

-Peter Singer and William MacAskill, April 2015

Part I

WHAT IS EFFECTIVE ALTRUISM?

1

The Drowning Child and the Expanding Circle

Peter Singer 1

To challenge my students to think about the ethics of what we owe to people in need, I ask them to imagine that their route to the university takes them past a shallow pond. One morning, I say to them, you notice a child has fallen in and appears to be drowning. To wade in and pull the child out would be easy but it will mean that you get your clothes wet and muddy, and by the time you go home and change you will have missed your first class.

I then ask the students: do you have any obligation to rescue the child? Unanimously, the students say they do. The importance of saving a child so far outweighs the cost of getting one's clothes muddy and missing a class, that they refuse to consider it any kind of excuse for not saving the child. Does it make a difference, I ask, that there are other people walking past the pond who would equally be able to rescue

¹Reprinted by kind permission of New Internationalist. Copyright New Internationalist. www.newint.org. Initially published in April 1997.

the child but are not doing so? No, the students reply, the fact that others are not doing what they ought to do is no reason why I should not do what I ought to do.

Once we are all clear about our obligations to rescue the drowning child in front of us, I ask: would it make any difference if the child were far away, in another country perhaps, but similarly in danger of death, and equally within your means to save, at no great cost — and absolutely no danger — to yourself? Virtually all agree that distance and nationality make no moral difference to the situation. I then point out that we are all in that situation of the person passing the shallow pond: we can all save lives of people, both children and adults, who would otherwise die, and we can do so at a very small cost to us: the cost of a new CD, a shirt or a night out at a restaurant or concert, can mean the difference between life and death to more than one person somewhere in the world — and overseas aid agencies like Oxfam overcome the problem of acting at a distance.

At this point the students raise various practical difficulties. Can we be sure that our donation will really get to the people who need it? Doesn't most aid get swallowed up in administrative costs, or waste, or downright corruption? Isn't the real problem the growing world population, and is there any point in saving lives until the problem has been solved? These questions can all be answered: but I also point out that even if a substantial proportion of our donations were wasted, the cost to us of making the donation is so small,

compared to the benefits that it provides when it, or some of it, does get through to those who need our help, that we would still be saving lives at a small cost to ourselves — even if aid organizations were much less efficient than they actually are.

I am always struck by how few students challenge the underlying ethics of the idea that we ought to save the lives of strangers when we can do so at relatively little cost to ourselves. At the end of the nineteenth century WH Lecky wrote of human concern as an expanding circle which begins with the individual, then embraces the family and 'soon the circle... includes first a class, then a nation, then a coalition of nations, then all humanity, and finally, its influence is felt in the dealings of man [sic] with the animal world'. On this basis the overwhelming majority of my students seem to be already in the penultimate stage — at least — of Lecky's expanding circle. There is, of course, for many students and for various reasons a gap between acknowledging what we ought to do, and doing it; but I shall come back to that issue shortly.

Our century is the first in which it has been possible to speak of global responsibility and a global community. For most of human history we could affect the people in our village, or perhaps in a large city, but even a powerful king could not conquer far beyond the borders of his kingdom. When Hadrian ruled the Roman Empire, his realm covered most of the "known" world, but today when I board a jet in London leaving what used to be one of the far-flung outposts of the Roman Empire, I pass over its opposite boundary before I am

even halfway to Singapore, let alone to my home in Australia. Moreover no matter what the extent of the empire, the time required for communications and transport meant that there was simply no way in which people could make any difference to the victims of floods, wars, or massacres taking place on the other side of the globe. By the time anyone had heard of the events and responded, the victims were dead or had survived without assistance. "Charity begins at home" made sense, because it was only "at home" — or at least in your own town — that you could be confident that your charity would make any difference.

Instant communications and jet transport have changed all that. A television audience of two billion people can now watch hungry children beg for food in an area struck by famine, or they can see refugees streaming across the border in search of a safe place away from those they fear will kill them. Most of that huge audience also have the means to help people they are seeing on their screens. Each one of us can pull out a credit card and phone in a donation to an aid organization which can, in a few days, fly in people who can begin distributing food and medical supplies. Collectively, it is also within the capacity of the United Nations — with the support of major powers — to put troops on the ground to protect those who are in danger of becoming victims of genocide.

Our capacity to affect what is happening, anywhere in the world, is one way in which we are living in an era of global responsibility. But there is also another way that offers an even

more dramatic contrast with the past. The atmosphere and the oceans seemed, until recently, to be elements of nature totally unaffected by the puny activities of human beings. Now we know that our use of chlorofluorocarbons has damaged the ozone shield; our emission of carbon dioxide is changing the climate of the entire planet in unpredictable ways and raising the level of the sea; and fishing fleets are scouring the oceans, depleting fish populations that once seemed limitless to a point from which they may never recover. In these ways the actions of consumers in Los Angeles can cause skin cancer among Australians, inundate the lands of peasants in Bangladesh, and force Thai villagers who could once earn a living by fishing to work in the factories of Bangkok.

In these circumstances the need for a global ethic is inescapable. Is it nevertheless a vain hope? Here are some reasons why it may not be.

We live in a time when many people experience their lives as empty and lacking in fulfilment. The decline of religion and the collapse of communism have left but the ideology of the free market whose only message is: consume, and work hard so you can earn money to consume more. Yet even those who do reasonably well in this race for material goods do not find that they are satisfied with their way of life. We now have good scientific evidence for what philosophers have said throughout the ages: once we have enough to satisfy our basic needs, gaining more wealth does not bring us more happiness.

Consider the life of Ivan Boesky, the multimillionaire Wall Street dealer who in 1986 pleaded guilty to insider trading. Why did Boesky get involved in criminal activities when he already had more money than he could ever spend? Six years after the insider-trading scandal broke, Boesky's estranged wife Seema spoke about her husband's motives in an interview with Barbara Walters for the American ABC Network's 20/20 program. Walters asked whether Boesky was a man who craved luxury. Seema Boesky thought not, pointing out that he worked around the clock, seven days a week, and never took a day off to enjoy his money. She then recalled that when in 1982 Forbes magazine first listed Boesky among the wealthiest people in the US, he was upset. She assumed he disliked the publicity and made some remark to that effect. Boesky replied: "That's not what's upsetting me. We're noone. We're nowhere. We're at the bottom of the list and I promise you I won't shame you like that again. We will not remain at the bottom of that list."

We must free ourselves from this absurd conception of success. Not only does it fail to bring happiness even to those who, like Boesky, do extraordinarily well in the competitive struggle; it also sets a social standard that is a recipe for global injustice and environmental disaster. We cannot continue to see our goal as acquiring more and more wealth, or as consuming more and more goodies, and leaving behind us an even larger heap of waste.

We tend to see ethics as opposed to self-interest; we assume that those who make fortunes from insider trading are successfully following self-interest — as long as they don't get caught — and ignoring ethics. We think that it is in our interest to take a more senior better-paid position with another company, even though it means that we are helping to manufacture or promote a product that does no good at all, or is environmentally damaging. On the other hand, those who pass up opportunities to rise in their career because of ethical "scruples" about the nature of the work, or who give away their wealth to good causes, are thought to be sacrificing their own interest in order to obey the dictates of ethics.

Many will say that it is naive to believe that people could shift from a life based on consumption, or on getting on top of the corporate ladder, to one that is more ethical in its fundamental direction. But such a shift would answer a palpable need. Today the assertion that life is meaningless no longer comes from existentialist philosophers who treat it as a shocking discovery: it comes from bored adolescents for whom it is a truism. Perhaps it is the central place of self-interest, and the way in which we conceive of our own interest, that is to blame here. The pursuit of self-interest, as standardly conceived, is a life without any meaning beyond our own pleasure or individual satisfaction. Such a life is often a self-defeating enterprise. The ancients knew of the "paradox of hedonism", according to which the more explicitly we pursue our desire for pleasure, the more elusive we will find its satisfaction. There is no reason to believe that human nature has changed so dramatically as to render the ancient wisdom inapplicable.

Here ethics offer a solution. An ethical life is one in which we identify ourselves with other, larger, goals, thereby giving meaning to our lives. The view that there is harmony between ethics and enlightened self-interest is an ancient one, now often scorned. Cynicism is more fashionable than idealism. But such hopes are not groundless, and there are substantial elements of truth in the ancient view that an ethically reflective life is also a good life for the person leading it. Never has it been so urgent that the reasons for accepting this view should be widely understood.

In a society in which the narrow pursuit of material self-interest is the norm, the shift to an ethical stance is more radical than many people realize. In comparison with the needs of people going short of food in Rwanda, the desire to sample the wines of Australia's best vineyards pales into insignificance. An ethical approach to life does not forbid having fun or enjoying food and wine; but it changes our sense of priorities. The effort and expense put into fashion, the endless search for more and more refined gastronomic pleasures, the added expense that marks out the luxury-car market — all these become disproportionate to people who can shift perspective long enough to put themselves in the position of others affected by their actions. If the circle of ethics really does expand, and a higher ethical consciousness spreads, it will fundamentally change the society in which we live.

2

What is Effective Altruism?

William MacAskill 1

Most forms of do-gooding start out with a What ("I want to promote microfinance!"), move to a How ("maybe I should do a sponsored marathon?") and simply take the Why for granted ("because of course microfinance is good!").

Effective altruism, in contrast, starts with a Why and a How, and lets them determine the What. Let me explain:

The Why is to make the world as good a place as it can possibly be. Rather than merely aiming to make the world better than when we found it — "to make a difference" — we want to make the most difference. So, for example, rather than simply trying to find a development charity that "does good work", Giving What We Can seeks to find those charities that do the very most to help people in developing countries with every pound or dollar they receive. In general, we seek out those activities that will do the most good with our time or money.

¹This article was initially posted to the Effective Altruism blog in May 2013 (www.effective-altruism.com/what-effective-altruism/).

The How — how to find those activities that do the most good — is by using robust evidence and careful reasoning. Where a question concerns a matter of fact, we try to find the best empirical evidence that is relevant to that question. (An anecdote is bad, a double-blind randomized controlled trial is better, a well-performed meta-analysis is best.) Where a question concerns values, we use clear arguments, rational reflection, and the latest insights from ethics, economics, and psychology to help us come to the right view. So, for example, rather than going with feel-good slogans like "follow your passion", or passing on anecdotes about specific people, at 80,000 Hours we're busy digging into all the available academic research related to doing good through your career, and clarifying, conceptually, what making a difference involves.

From these two ideas, the What follows. Effective altruists currently tend to think that the most important causes to focus on are global poverty, factory farming, and the long-term future of life on Earth. In each case, the reasoning is that the stakes are very high, and there is the potential to make a lot of progress. Right now, within the Centre for Effective Altruism, the What consists of organizations that promote donating a good chunk of your income to the world's most effective charities (Giving What We Can and The Life You Can Save); or that advise individuals on which careers enable them to have the greatest positive impact (80,000 Hours); or that try to figure out how best to improve animal welfare (Animal Charity Evaluators). But these activities are just our

current best guesses. If we had good evidence or arguments that showed that we could do more good by doing something else, then we'd do that instead.

3

Scope Neglect

Eliezer Yudkowsky 1

Imagine that 2,000, 20,000 or 200,000 migrating birds die each year by drowning in uncovered oil ponds, which the birds mistake for bodies of water. These deaths could be prevented by covering the oil ponds with nets. How much money would you be willing to pay to provide the needed nets?

Three groups of subjects considered three versions of the above question, asking them how high a tax increase they would accept to save 2,000, 20,000, or 200,000 birds. The response, known as stated willingness to pay, or SWTP, had a mean of \$80 for the 2,000-bird group, \$78 for 20,000 birds, and \$88 for 200,000 birds (Desvousges et al).² This phenomenon is known as scope insensitivity or scope neglect.

¹This article is excerpted from the 2008 volume Global Catastrophic Risks by Nick Bostrom and Milan Cirkovic and the blog post Scope Sensitivity by Eliezer Yudkowsky in 2007, www.lesswrong.com/lw/hw/scope_insensitivity.

²Desvousges, William H., F. Reed Johnson, Richard W. Dunford, Kevin J. Boyle, Sara P. Hudson, and K. Nicole Wilson. 1993. "Measuring Natural Resource Damages with Contingent Valuation: Tests of Validity and Reliability." In Hausman 1993, 91–164.

Similar studies have shown that Toronto residents would pay little more to clean up all polluted lakes in Ontario than polluted lakes in a particular region of Ontario,³ and that residents of four western US states would pay only 28% more to protect all 57 wilderness areas in those states than to protect a single area.⁴.

The most widely accepted explanation for scope neglect appeals to the affect heuristic. Kahneman, Ritov, and Schkade write:

The story constructed by Desvousges et al probably evokes for many readers a mental representation of a prototypical incident, perhaps an image of an exhausted bird, its feathers soaked in black oil, unable to escape. The hypothesis of valuation by prototype asserts that the affective value of this image will dominate expressions of the attitude to the problem — including the willingness to pay for a solution. Valuation by prototype implies extension neglect.⁵

³Kahneman, Daniel. "Comments by Professor Daniel Kahneman." In Valuing Environmental Goods: An Assessment of the Contingent Valuation Method, edited by Ronald G. Cummings, David S. Brookshire, and William D. Schulze, (1986), 226–235. Vol. 1.B. Experimental Methods for Assessing Environmental Benefits. Totowa, NJ: Rowman & Allanheld. http://yosemite.epa.gov/ee/epa/eerm.nsf/vwAN/EE-0280B-04.pdf/\$file/EE-0280B-04.pdf.

⁴McFadden, Daniel L., and Gregory K. Leonard. "Issues in the Contingent Valuation of Environmental Goods: Methodologies for Data Collection and Analysis." In Hausman 1993, 165–215.

⁵Kahneman, Daniel, Ilana Ritov, and David Schkade. "Economic Preferences or Attitude Expressions?: An Analysis of Dollar Responses to Public Issues." *Journal of Risk and Uncertainty* 19.1–3 (1999):203–235.

Two other hypotheses accounting for scope neglect include purchase of moral satisfaction⁶ and good cause dump.⁷

Purchase of moral satisfaction suggests that people spend enough money to create a "warm glow" in themselves, and the amount required is a property of the person's psychology, having nothing to do with birds. Good cause dump suggests that people have some amount of money they are willing to pay for "the environment," and any question about environmental goods elicits this amount.

Scope neglect has been shown to apply to human lives. Carson and Mitchell (1995) report that increasing the alleged risk associated with chlorinated drinking water from 0.004 to 2.43 annual deaths per 1,000 (a factor of 600) increased SWTP from \$3.78 to \$15.23 (a factor of 4). Baron and Greene found no effect from varying lives saved by a factor of ten.⁸

Fetherstonhaugh et al, in a paper entitled "Insensitivity to the Value of Human Life: A Study of Psychophysical Numbing," found evidence that our perception of human deaths, and valuation of human lives, obeys Weber's Law — meaning

⁶Kahneman, Daniel, and Jack L. Knetsch. "Valuing Public Goods: The Purchase of Moral Satisfaction." *Journal of Environmental Economics and Management* 22.1 (1992), 57–70.

⁷Carson, Richard T., and Robert Cameron Mitchell. 1995. "Sequencing and Nesting in Contingent Valuation Surveys." *Journal of Environmental Economics and Management* 28.2 (1992), 155–173.

⁸Baron, Jonathan, and Joshua D. Greene. 1996. "Determinants of Insensitivity to Quantity in Valuation of Public Goods: Contribution, Warm Glow, Budget Constraints, Availability, and Prominence." *Journal of Experimental Psychology: Applied 2.2* (1996), 107–125.

that we use a logarithmic scale. And indeed, studies of scope neglect in which the quantitative variations are huge enough to elicit any sensitivity at all, show small linear increases in willingness to pay corresponding to exponential increases in scope. Kahneman, Ritov, and Schkade interpret this as an additive effect of scope affect and prototype affect — the prototype image elicits most of the emotion, and the scope elicits a smaller amount of emotion which is added (not multiplied) with the first amount.

Albert Szent-Györgyi said: "I am deeply moved if I see one man suffering and would risk my life for him. Then I talk impersonally about the possible pulverization of our big cities, with a hundred million dead. I am unable to multiply one man's suffering by a hundred million." Human emotions take place within an analog brain. The human brain cannot release enough neurotransmitters to feel emotion a thousand times as strong as the grief of one funeral. A prospective risk going from 10,000,000 deaths to 100,000,000 deaths does not multiply by ten the strength of our determination to stop it. It adds one more zero on paper for our eyes to glaze over, an effect so small that we must usually jump several orders of magnitude to detect the difference experimentally.

⁹Fetherstonhaugh, David, Paul Slovic, Stephen M. Johnson, and James Friedrich. 1997. "Insensitivity to the Value of Human Life: A Study of Psychophysical Numbing." *Journal of Risk and Uncertainty* 14.3 (1997), 283–300.

The moral is that if you want to be an effective altruist, you have to think it through with the part of your brain that processes those unexciting inky zeroes on paper, not just the part that gets real worked up about that poor struggling oilsoaked bird.

4

Tradeoffs

Julia Wise 1

Economists love to think about tradeoffs (or opportunity costs, as they call them). Any money we spend can't be spent on something else, so if I use \$2.50 to buy a strawberry milk-shake it means I'm not using that \$2.50 to get the chocolate one or the mint chocolate chip one.

That's pretty easy to think about. But it also means I'm not getting a bus fare, a light bulb, or anything else with that \$2.50. And if I buy that strawberry milkshake, according to standard economics it means there's nothing else in the world I would rather buy with that money.

I don't think we're usually that rational.

For one thing, it's unpleasant to think about negatives. We like to think about what our money *does* get us, rather than the infinite variety of things it *doesn't* get. Also, there are so many alternatives that we can't really consider them all every time we spend money.

¹This article was initially posted to Julia Wise's blog in March 2012 at http://www.givinggladly.com/2012/03/tradeoffs.html.

I once saw a flippant proposal that we draw people's attention to this in a gruesome way by labeling all prices in Dead Child Currency². If it costs \$800 to save a child's life, each \$800 spent on anything else ... you get the idea.

I used to make myself think that way. Before I parted with any money, I'd ask myself what it could do for a woman in Africa who needed the money to take care of her family. (It doesn't have to be her, but that's who I always imagined.) Did I value my new shoes more than her month's groceries? More than her children's vaccinations or school fees? Could I make that tradeoff?

Sometimes I made it and felt awful afterwards. Eventually, my husband and I decided to draw a firm line about what money was to be donated and what was ours. Each year, we decided what chunk of our budget would to to the best charities we could find. The rest was ours to spend on anything else we wanted. We've edged the charity budget upward over the years, and now it's at 50% of our income. Making that decision once a year is a lot easier than making it every time you stand in a checkout line.

I recently met a young man who was seriously thinking these things over. "But isn't it right to think about the trade-offs?" he asked. "Shouldn't we be thinking about them all the time?"

 $^{^2{\}rm This}$ satirical proposal was described in Dead Children Currency by Scott Alexander, at www.raikoth.net/deadchild.html

I think it's good to go through a period of thinking that way. When you live in another country, you become able to understand prices without converting back to your own currency. In the same way, when you start thinking about all your spending in Vaccination Currency or Mosquito Net Currency it becomes habitual. Your spending can't help but be affected.

I also think there's only so much grief we can carry. I cannot spend the next 60 years counting dead children on every receipt. I would break.

So my advice is to spend a while really noticing that tradeoff. Notice whether you really do value the milkshake more than a child's vaccination. Think about what choice you would want to be made if it were your child that needed help. And then, after a while, make yourself a budget that reflects those values. Set aside some money for unnecessary things that make you happy. Do what you think will nurture you to age 90 as a generous and strategic giver. Because that, in the end, is what will help the most people.

Part II CHARITY EVALUATION

5

Efficient Charity: Do Unto Others

Scott Alexander 1

Imagine you are setting out on a dangerous expedition through the Arctic on a limited budget. The grizzled old prospector at the general store shakes his head sadly: you can't afford everything you need; you'll just have to purchase the bare essentials and hope you get lucky. But what is essential? Should you buy the warmest parka, if it means you can't afford a sleeping bag? Should you bring an extra week's food, just in case, even if it means going without a rifle? Or can you buy the rifle, leave the food, and hunt for your dinner?

And how about the field guide to Arctic flowers? You like flowers, and you'd hate to feel like you're failing to appreciate the harsh yet delicate environment around you. And a digital camera, of course — if you make it back alive, you'll have to put the Arctic expedition pics up on Facebook. And a hand-

¹This article was initially posted on LessWrong in December 2010 at www.lesswrong.com/lw/3gj/efficient_charity_do_unto_others/.

crafted scarf with authentic Inuit tribal patterns woven from organic fibers! Wicked!

...But of course buying any of those items would be insane. The problem is what economists call opportunity costs: buying one thing costs money that could be used to buy others. A hand-crafted designer scarf might have some value in the Arctic, but it would cost so much it would prevent you from buying much more important things. And when your life is on the line, things like impressing your friends and buying organic pale in comparison. You have one goal — staying alive — and your only problem is how to distribute your resources to keep your chances as high as possible. These sorts of economics concepts are natural enough when faced with a journey through the freezing tundra.

But they are decidedly not natural when facing a decision about charitable giving. Most donors say they want to "help people". If that's true, they should try to distribute their resources to help people as much as possible. Most people don't. In the "Buy A Brushstroke" campaign, eleven thousand British donors gave a total of £550,000 to keep the famous painting "Blue Rigi" in a UK museum. If they had given that £550,000 to buy better sanitation systems in African villages instead, the latest statistics suggest it would have saved the lives of about three hundred people from disease. Each individual £50 donation could have given a year of normal life back to a Third Worlder afflicted with a disabling condition like blindness or limb deformity.

Most of those 11,000 donors genuinely wanted to help people by preserving access to the original canvas of a beautiful painting. And most of those 11,000 donors, if you asked, would say that a thousand people's lives are more important than a beautiful painting, original or no. But these people didn't have the proper mental habits to realize that was the choice before them, and so a beautiful painting remains in a British museum and somewhere in the Third World a thousand people are dead.

If you are to "love your neighbor as yourself", then you should be as careful in maximizing the benefit to others when donating to charity as you would be in maximizing the benefit to yourself when choosing purchases for a polar trek. And if you wouldn't buy a pretty picture to hang on your sled in preference to a parka, you should consider not helping save a famous painting in preference to helping save a thousand lives.

Not all charitable choices are as simple as that one, but many charitable choices do have right answers. GiveWell.org, a site which collects and interprets data on the effectiveness of charities, predicts that antimalarial drugs save one child from malaria per £5,000 worth of medicine, but insecticide-treated bed nets save one child from malaria per £500 worth of netting. If you want to save children, donating bed nets instead of antimalarial drugs is the objectively right answer, in the same way as buying a £500 TV instead of an identical TV that costs £5,000 is the right answer. And since saving a child from diarrheal disease costs £5,000, donating to an organization

fighting malaria instead of an organization fighting diarrhea is the right answer, unless you are donating based on some criteria other than whether you're helping children or not.

Say all of the best Arctic explorers agree that the three most important things for surviving in the Arctic are good boots, a good coat, and good food. Perhaps they have run highly unethical studies in which they release thousands of people into the Arctic with different combination of gear, and consistently find that only the ones with good boots, coats, and food survive. Then there is only one best answer to the question "What gear do I buy if I want to survive?" — good boots, good food, and a good coat. Your preferences are irrelevant; you may choose to go with alternate gear, but only if you don't mind dying.

And likewise, there is only one best charity: the one that helps the most people the greatest amount per dollar. This is vague, and it is up to you to decide whether a charity that raises forty schoolchildren's marks by one letter grade for £100 helps people more or less than one that prevents one fatal case of tuberculosis per £100 or one that saves twenty acres of rainforest per £100. But you cannot abdicate the decision, or you risk ending up like the 11,000 people who accidentally decided that a pretty picture was worth more than a thousand people's lives.

Deciding which charity is the best is hard. It may be straightforward to say that one form of antimalarial therapy is more effective than another. But how do both compare to financing medical research that might or might not develop a "magic bullet" cure for malaria? Or financing development of a new kind of supercomputer that might speed up all medical research? There is no easy answer, but the question has to be asked.

What about just comparing charities on overhead costs, the one easy-to-find statistic that's universally applicable across all organizations? This solution is simple, elegant, and wrong. High overhead costs are only one possible failure mode for a charity. Consider again the Arctic explorer, trying to decide between a £200 parka and a £200 digital camera. Perhaps a parka only cost £100 to make and the manufacturer takes £100 profit, but the camera cost £200 to make and the manufacturer is selling it at cost. This speaks in favor of the moral qualities of the camera manufacturer, but given the choice the explorer should still buy the parka does something vital inefficiently. A parka sold at cost would be best, but in its absence the explorer shouldn't hesitate to choose the parka over the camera.

The same applies to charity. An antimalarial net charity that saves one life per £500 with 50% overhead is better than an antidiarrheal drug charity that saves one life per £5000 with 0% overhead: £10,000 donated to the high-overhead charity will save ten lives; £10,000 to the lower-overhead charity will only save two. Here the right answer is to donate to the antimalarial charity while encouraging it to find ways to lower its overhead. In any case, examining the financial practices of

a charity is helpful but not enough to answer the "which is the best charity?" question.

Just as there is only one best charity, there is only one best way to support that charity. Whether you volunteer or donate money or raise awareness is your own choice, but that choice has consequences. If a high-powered lawyer who makes £1,000 an hour chooses to take an hour off to help clean up litter on the beach, he's wasted the opportunity to work the full day, make the extra £1,000, donate to a charity that will hire a hundred poor people for £10/hour to clean up litter, and end up with a hundred times more litter removed. If he went to the beach because he wanted the sunlight and the fresh air and the warm feeling of personally contributing to something, that's fine. If he actually wanted to help people by beautifying the beach, he's chosen an objectively wrong way to go about it. And if he wanted to help people, period, he's chosen a very wrong way to go about it, since that £1,000 could save two people from malaria. Unless the litter he removed is really worth more than two people's lives to him, he's erring even according to his own value system.

...And the same is true if his philanthropy leads him to work full-time at a non-profit instead of going to law school to become a lawyer who makes £1,000 / hour in the first place. Unless it's one *hell* of a non-profit.

The Roman historian Sallust said of Cato "He preferred to be good, rather than to seem so". The lawyer who quits a high-powered law firm to work at a non-profit organization certainly seems like a good person. But if we define "good" as helping people, then the lawyer who stays at his law firm but donates the profit to charity is taking Cato's path of maximizing how much good he does, rather than how good he looks.

And this dichotomy between being and seeming good applies not only to looking good to others, but to ourselves. When we donate to charity, one incentive is the warm glow of a job well done. A lawyer who spends his day picking up litter will feel a sense of personal connection to his sacrifice and relive the memory of how nice he is every time he and his friends return to that beach. A lawyer who works overtime and donates the money online to starving orphans in Romania may never get that same warm glow. But concern with a warm glow is, at root, concern about seeming good rather than being good — albeit seeming good to yourself rather than to others. There's nothing wrong with donating to charity as a form of entertainment if it's what you want — giving money to the Art Fund may well be a quicker way to give yourself a warm feeling than seeing a romantic comedy at the cinema — but charity given by people who genuinely want to be good and not just to feel that way requires more forethought.

It is important to be rational about charity for the same reason it is important to be rational about Arctic exploration: it requires the same awareness of opportunity costs and the same hard-headed commitment to investigating efficient use of resources, and it may well be a matter of life and death.

6

"Efficiency" Measures Miss the Point

Dan Pallotta ¹

An e-mail I got from a former employee last week exemplified a dangerous public mythology: "You see, for every dollar a donor gives they have the expectation that it's used efficiently. After all they have choices, they can give that same dollar to another charity. Donors want their donations to go as far as possible..."

There are two fatal errors here. The first is that high administrative efficiency equals high impact. It doesn't. The second is that the admin-to-program ratio is measuring efficiency. If it isn't measuring impact, it's axiomatic that it isn't measuring efficiency, because the only efficiency that matters is the efficiency associated with impact. Take the frugal breast cancer charity that consistently fails to find a cure for breast cancer. The last word a woman dying of breast cancer would

¹This article was initially posed to Dan Pallotta's blog in June 2009 at http://blogs.hbr.org/2009/06/efficiency-measures-miss-the-p.

use to describe it would be "efficient." Not if she factors in the value of her life.

As for making donations "go as far as possible," consider two soup kitchens. Soup kitchen A reports that 90% of every donation goes to the cause. Soup Kitchen B reports 70%. You should donate to A, right? No-brainer. Unless you actually visited the two and found that the so-called more "efficient" Soup Kitchen A serves rancid soup in a dilapidated building with an unpleasant staff and is closed half the time, while Soup Kitchen B is open 24/7, and has a super-friendly staff that serves nutritious soup in a state-of-the-art facility. Now which looks better?

The admin-to-program ratio would have failed you completely. It betrays your trust. It's utterly deficient in data about which soup kitchen is better at serving soup. It undermines your compassion and insults your contribution. And yet we praise it as a yardstick of morality and trustworthiness. It's the exact opposite.

We should stop saying charities with low ratios are efficient. Efficient at what? Fundraising? "Inefficient" — as in expedient — fundraising may accelerate problem-solving, making its "inefficiency" efficient in the big picture. Say Jonas Salk spent \$50 million to raise \$100 million to find a polio vaccine. The admin-to-program ratio would report he had a shameful 50% overhead. But the \$100 million he raised wasn't his end result. His end result was a vaccine. Divide the \$50 million fundraising expense into the God-only-knows-

how-many billions of dollars a polio vaccine is worth, and his overhead ratio at eradicating polio is 0%. A hypothetical competing charity with 10% fundraising cost that comes up empty on a vaccine has 100% overhead against the goal of a vaccine, because it never found one. But it's labeled the more "efficient" charity. As one of millions who dodged polio because of Salk, I'd have to disagree.

Let's get unhypothetical. In 1995, Physicians for Human Rights had revenues of approximately \$1.3 million. They spent approximately \$750,000, or 58 percent of revenues, on programs. Today that organization would fail all of the watchdog standards for "efficiency". It would be ineligible for a BBB Wise Giving Alliance seal of approval. The Nobel Peace Prize committee felt differently. Physicians for Human Rights won the Nobel Prize in 1997 for its work as a founding member of the International Campaign to Ban Landmines. Imagine coming out of a shoe store with a brand new pair of shoes full of holes, and whispering to your friends, "You wouldn't believe how low the overhead was on these shoes." That's exactly what Americans are doing with hundreds of billions of annual charitable donations. We take huge pride in giving to charities with low overhead without knowing a damned thing about whether they're any good at what they do.

The e-mail from my former colleague was right in one respect. Donors do have a choice. And they should stop using this hallucinogenic "efficiency" ratio to determine how they make it.

7

How Not to Be a "White in Shining Armor"

Holden Karnofsky 1

GiveWell's current top-rated charities focus on proven, costeffective health interventions. These interventions appear to solve certain problems (malaria, parasites) quite well, while making no direct attempt to solve other problems (economic growth, education, gender equity, and more). One of the common lines of objection we get to these recommendations goes something like: "Why should I put all my money into fighting malaria, ignoring other important problems? Isn't it unethical to ignore the other essential needs?"

We believe this objection commits the common fallacy of viewing the developed-world donor as the only person who can improve things for the beneficiaries: one term for taking this mentality too far is "Whites in Shining Armor". Often, in the media and in communications from non-profits, global

¹This article was initially posted to the GiveWell blog in April 2012 at http://blog.givewell.org/2012/04/12/how-not-to-be-a-white-in-shining-armor.

poverty is presented as a simple fight between local problems and developed-world heroes. The problem is that as outsiders, we often have very poor understanding of the true dynamics behind overseas problems — and by attempting to solve problems that we understand poorly, we can make things worse.

We fundamentally believe that progress on most problems must be locally driven. So we seek to improve people's abilities to make progress on their own, rather than taking personal responsibility for each of their challenges. The question is, how can we best accomplish this?

Locally driven projects

A common and intuitively appealing answer is to let locals drive philanthropic projects. This answer has some appeal for us; we have previously written about, and given a small amount of money to, "low-insulation charities" that seem adaptive, locally connected, and overall driven by local needs rather than donors' plans.

At the same time, we have noted some major challenges caused by doing things this way. Which locals should be put in charge? There is the inherent risk that the people who least need help will be best positioned to get involved with making the key decisions: on our visit to India, we noted that some of the organizations in charge seemed to consist simply of local elites making ad-hoc decisions. To truly reach those who most

need help seems to require being "systematically bottom-up," a more complex and difficult approach.

Global health and nutrition

Another approach to "putting locals in the driver's seat" is quite different. It comes down to acknowledging that as funders, we will always be outsiders, so we should focus on helping with what we're good at helping with and leave the rest up to locals.

Here I think an analogy to helping friends and family is somewhat illustrative. I try to help my friends and family in domains in which I'm relatively knowledgeable (for example, computer issues) and I tend not to put much effort into helping in other areas in which I'm not so knowledgeable (for example, picking clothes) even if the latter are more important issues for them. I know I appreciate when my friends and family deal with me this way, and I don't appreciate people who are determined to help me in domains that they don't understand very well (even if these domains are very important to me).

We believe that the track record of outside aid points to health and nutrition as the areas that developed-world outsiders understand best and are best-positioned to help with.

It's not that we think global health and nutrition are the only important, or even the most important, problems in the developing world. It's that we're trying to focus on what we can do well, and thus maximally empower people to make locally-driven progress on other fronts.

Cash transfers

One more approach to "putting locals in the driver's seat" is to ensure that money goes directly into the pockets of local residents, to do with as they see fit. This is the work done by the charity GiveDirectly, which supports unconditional cash transfers. For those who are even more concerned than we are about the trap of "whites in shining armor," this option has promise.

8

Estimation Is the Best We Have

Katja Grace 1

This argument seems common to many debates:

Proposal P arrogantly assumes that it is possible to measure X, when really X is hard to measure and perhaps even changes depending on other factors. Therefore we shouldn't do P.

This could make sense if X wasn't especially integral to the goal. For instance if the proposal were to measure short distances by triangulation with nearby objects, a reasonable criticism would be that the angles are hard to measure, relative to measuring the distance directly. But this argument is commonly used in situations where optimizing X is the whole point of the activity, or a large part of it.

¹This article was initially posted to Katja Grace's blog in December 2010 at www.meteuphoric.wordpress.com/2010/12/20/estimation-is-the-best-we-have.

Criticism of cost-benefit approaches to doing good provides a prime example. A common argument is that it's just not possible to tell if you are increasing net welfare, or by how much. The critic concludes then that a different strategy is better, for instance some sort of intuitive adherence to strict behavioral rules.

But if what we fundamentally think matters most is increasing welfare, or at least reducing extreme suffering, then the difficulty of doing the associated mathematics perfectly should not warrant abandoning the goal. Better to put your effort into increasing your accuracy than to throw it away on a strategy that is more random with regard to your goal.

A CEO would sound ridiculous making this argument to his shareholders: 'You guys are being ridiculous. It's just not possible to know which actions will increase the value of the company by exactly how much. Why don't we try to make sure that all of our meetings end on time instead?'

In general, when optimizing X is integral to the goal, the argument must fail. If for instance the point is to make X as close to three as possible, then no matter how bad your best estimates are of what X will be under different conditions, you can't do better by ignoring X altogether. If you have a non-estimating-X strategy which you anticipate would do better than your best estimate in getting a good value of X, then you in fact believe yourself to have a better estimating-X strategy.

Probabilistic risk assessment is claimed by some to be impossibly difficult. People are often wrong, and may fail to think of certain contingencies in advance. So if for example we want to know how prepared to be for a nuclear war, we should analyze relevant scenarios quantitatively. This could be a defensible position. Perhaps intuitions can better assess probabilities implicitly via some other activity than explicitly thinking about them.

However I have not heard this claim accompanied by any motivating evidence. And even if it were true, it would likely make sense to convert the qualitative assessments into quantitative ones and aggregate them with information from other sources, rather than disregarding quantitative assessments altogether.

9

Our Updated Top Charities

Elie Hassenfeld 1

Our top charities are (in alphabetical order):

- Against Malaria Foundation (AMF)
- Deworm the World Initiative (DtWI), led by Evidence Action
- GiveDirectly
- Schistosomiasis Control Initiative (SCI)

We have also included four additional organizations on our top charities page as standout charities. They are (in alphabetical order):

- Development Media International (DMI)
- The Global Alliance for Improved Nutrition's Universal Salt Iodization program (GAIN-USI)

 $^{^1\}mathrm{This}$ article was initially posted to the GiveWell Blog in December 2014 at blog.givewell.org/2014/12/01/our-updated-top-charities/.

- International Council for the Control of Iodine Deficiency Disorders Global Network (ICCIDD)
- · Living Goods

We see a strong possibility that these will become top-tier recommended charities in the future, and we can see reasons that impact-minded donors could choose to support them today.

Against Malaria Foundation $(AMF)^2$

Important changes in the last 12 months

We named AMF our #1-ranked charity at the end of 2011. Over the next 2 years, AMF received more than \$10 million on the basis of our recommendation but struggled to identify opportunities to use the funds it had received. At the end of 2013, we announced that we planned not to recommend additional donations to AMF until it committed the bulk of its current funds. This did not reflect a negative view of AMF; instead it reflected room for more funding related issues.

In 2014, AMF finalized several distributions in Malawi and the Democratic Republic of the Congo (DRC) with three different implementing partners (two of which account for the bulk of the nets to be distributed). In 2014, it committed approximately \$8.4 million to distributions which will take place

²GiveWell's full review of AMF is available at www.givewell.org/international/top-charities/AMF.

before January 1, 2016 (some of which have already begun) and now has \$6.8 million available for future distributions. \$1.7 million of this is committed to a distribution scheduled for 2017 (and could potentially be allocated to distributions taking place sooner). Excluding the 2017 distribution, AMF has committed approximately \$11.2 million to distributions in its history.

Funding gap

AMF requires access to funding in order to negotiate deals because it cannot initiate discussions with potential partners unless it is confident that it will have sufficient funding to support its future agreements. The funding it currently holds would enable it to fund approximately 3 distributions at a scale similar to what it has funded recently.

AMF has told us that it has a pipeline of possible future net distributions that add up to \$36 million. We see some reason for caution in thinking about AMF's room for more funding. It has made strong progress on being able to negotiate distributions and commit funds. However, as of today there have only been two large-scale distributions that have moved forward far enough for data to be available. Both of these are significantly smaller than distributions AMF has recently or will soon fund, and both are in the same area with the same partner as each other. Some of the recently negotiated distributions could prove more challenging (since they are in DRC).

If AMF received an additional \$10 million in total over the next 4 months, it would have about twice as much funding available as the total it committed to large-scale distributions in 2014. (As stated above, it committed \$8.4 million to distributions taking place before 2017 and has \$6.8 million available for further commitments.) If it received \$25 million, it would have about 4 times that total. 2–4 times past distributions seems like a range that would allow AMF to do significantly more than it has in the past, without going so far beyond its past capacity as to raise serious scaling concerns.

We believe that \$10 million total (the low end of that range), which means \$5 million after a grant by Good Ventures, is an appropriate target after which further donations are likely better off going to other charities.

Key considerations

Program impact and cost-effectiveness.

Our best guess is that distributing bednets is in the same costeffectiveness range as deworming programs and more costeffective than cash transfers by a factor of 5–10, although this is subject to substantial uncertainty.³

³GiveWell's cost-effectiveness analyses for bednets and other interventions are available at www.givewell.org/international/technical/criteria/costeffectiveness/cost-effectiveness-models

Directness and robustness of the case for impact.

We believe that the connection between AMF receiving funds and those funds helping very poor individuals is less direct than GiveDirectly's and more direct than SCI's or DtWI's. The uncertainty of our estimates is driven by a combination of AMF's challenges historically disbursing the funds it receives and a general recognition that aid programs, even those as straightforward as bednets, carry significant risks of failure via ineffective use of nets, insecticide resistance or other risks we don't yet recognize relative to GiveDirectly's program. AMF conducts extensive monitoring of its program; these results have generally indicated that people use the nets they receive.

Transparency and communication.

AMF has been extremely communicative and open with us. We feel we have a better understanding of AMF than SCI and worse than GiveDirectly. In particular, were something to go wrong in one of AMF's distributions, we believe we would eventually find out (something we are not sure of in the case of SCI), but we believe our understanding would be less quick and complete than it would be for problems associated with GiveDirectly's program (which has more of a track record of consistent intensive followup).

Risks

- Two of AMF's recent distributions (and much of its future pipeline) will take place in the DRC. Our impression is that the DRC is a particularly difficult place to work, and it is possible that AMF's distributions there will struggle or fail. We view this as a moderate risk.
- We are not highly confident that AMF will be able to finalize additional distributions and do so quickly. AMF could struggle again to agree to distribution deals, leading to long delays before it spends funds. We view this as a relatively minor risk because the likely worst case scenario is that AMF spends the funds slowly (or returns funds to donors).
- We remain concerned about the possibility of resistance to the insecticides used in bednets. There don't appear to be major updates on this front since our 2012 investigation into the matter; we take the lack of major news as a minor positive update.

A note on how quickly we expect AMF to spend the funds it receives. AMF works by sourcing, evaluating and negotiating deals for net distributions. This process takes time and requires AMF to have significant access to funding — it cannot approach a country to begin negotiations unless it is confident that it will have sufficient funding to pay for the nets it offers. We would not be surprised if AMF fails to reach additional

deals in the next 12 months. We do expect it to commit the majority of its available funds (that it will have as of this coming January) within the next 24 months and if not much progress in committing funds is made in the next 12 months, we will adjust our recommendation for 2015 accordingly.

Deworm the World Initiative, (DtWI), led by Evidence Action⁴

Important changes in the last 12 months

Dr. Kevin Croke released a new study of a randomized controlled trial of a deworming program showing large, long-term impacts from deworming programs. This study is a significant positive update on the impacts of deworming and increased our confidence that deworming programs have significant long-term impacts.

DtWI spent the funds it received due to GiveWell's recommendation largely as we anticipated; it now has some (though limited) room for more funding.

In 2014, two events affected DtWI's projection of the additional funding it would require to scale up in India:

• The Children's Investment Fund Foundation (CIFF), a major foundation that had supported DtWI's programs in Kenya, agreed to a 6-year, \$17.7 million grant to support DtWI's expansion to additional states in India

⁴GiveWell's full review of DtWI is available at www.givewell.org/international/top-charities/deworm-world-initiative.

and technical assistance to the Government of India for a national deworming program. With these funds, DtWI does not require significant additional funding to support its India expansion.

• The new Indian government expressed interest in conducting a single deworming day nationally with increased national attention and resources. Advocating for such a policy and assisting the national government in creating a plan became the major focus of DtWI's India work in 2014, which both reduced the amount of time it was able to spend generating interest in heavy DtWI involvement in new states and also required little funding since there were few costs of that project aside from staff time. We see this as positive news regarding DtWI's potential impact; it may simply reduce DtWI's further need for funds from individual donors.

Together, these changes led DtWI to the conclusion that funding is no longer the bottleneck to reaching more people in India.

Funding gap

DtWI told us that it seeks \$1.3 million over the next two years. We expect it to allocate approximately 30% of the additional funds it receives for work related to expanding school-based, mass deworming programs (including related operating and impact evaluation expenses) and will allocate other funds to priorities that are less directly connected to expanding and

evaluating deworming programs (investigating ways to combine other evidence-based programs with deworming rollouts, supplementing a project supported by another funder).

Good Ventures has announced a \$250,000 grant to DtWI, leaving it with \$1.05 million in remaining room for more funding over the next two years. We would ideally like DtWI to receive an additional \$500,000 (for a total of \$750,000) to provide it with more than half of its two-year gap.

Key considerations

Program impact and cost-effectiveness.

Our current calculations indicate that DtWI-associated deworming, when accounting for DtWI's potential "leverage" in influencing government funds, has extremely strong cost-effectiveness, better than bednets and 10–20 times better than cash transfers.⁵

Directness and robustness of the case for impact.

DtWI doesn't carry out deworming programs itself; it advocates for and provides technical assistance to governments implementing deworming programs, making direct assessments of its impact challenging. There are substantial potential advantages to supporting such an organization, as it may be able to have more impact per dollar by influencing government policy than by simply carrying out programs on its own, but

⁵GiveWell's cost-effectiveness estimates for deworming are available at www.givewell.org/international/technical/criteria/cost-effectiveness/costeffectiveness-models

this situation also complicates impact assessment. While we believe DtWI is impactful, our evidence is limited, and in addition, there is always a risk that future expansions will prove more difficult than past ones. In addition, DtWI is now largely raising funds to support research projects that are not directly connected to short-term implementation of deworming programs. We do not have a view about the value of these research projects.

Transparency and communication.

DtWI has been communicative and open with us. We have only recommended DtWI for one year and therefore have less history with it than AMF, GiveDirectly, or SCI, but we believe that were something to go wrong with DtWI's work, we would be able to learn about it and report on it.

Risks

- DtWI is part of a larger organization, Evidence Action, so changes that affect Evidence Action (and its other programs) could indirectly impact DtWI. For example, if a major event occurs (either positive or negative) for Evidence Action, it is likely that it would reduce the time some staff could devote to DtWI.
- Most of DtWI's funding is in the form of restricted funding from large, institutional funders. We are not sure how DtWI's plans would change in response to a large funder offering it significant support to undertake a project not directly in line with its current plans.

GiveDirectly⁶

Important changes in the last 12 months

GiveDirectly continued to scale up significantly, utilizing most of the funding it received at the end of last year. It continued to share informative and detailed monitoring information with us. Overall, it grew its operations while maintaining high quality.

In June, three of its board members launched Segovia, a for-profit company aimed at improving the efficiency of cash transfer distributions in the developing world.

GiveDirectly is working with other researchers to begin a very large study on cash transfers and the impact they have on broader economic factors such as inflation and job growth. This study will include a long-term follow up component as well. GiveDirectly told us that the ideal sample size for this study, which is randomized at the village level, would require \$15 million for cash transfers. Baseline data collection for the study began in August 2014. GiveDirectly has preregistered its plans for measurement and analysis.

Funding gap

GiveDirectly has scaled up significantly over the past year, spending (or committing to spend by enrolling recipients) approximately \$13.6 million of the \$17.4 million it received

⁶GiveWell's full review of GiveDirectly is available at www.givewell.org/international/top-charities/give-directly

last year. (It also allocated an additional \$1.8 million to other organizational costs.) It now believes that it could spend up to \$40 million in a year.

We believe this is a reasonable cap for GiveDirectly and would not hesitate to see it receive this amount. However, due to other charities' significantly superior estimated cost-effectiveness, we are seeking larger total amounts for them. We hope that GiveDirectly will receive at least \$1 million from individual donors (excluding Good Ventures) this giving season as a result of our recommendation.

Key considerations

Program impact and cost-effectiveness.

Our best guess is that deworming or distributing bednets achieves 5–10 times more humanitarian benefit per dollar donated than cash transfers.⁷

Directness and robustness of the case for impact.

GiveDirectly collects and shares a significant amount of relevant information about its activities. The data it collects show that it successfully directs cash to very poor people, that recipients generally spend funds productively (sometimes on food, clothing, or school fees, other times on investments in a business or home infrastructure), and that it leads to very low levels of interpersonal conflict and tension. We are more

⁷GiveWell's cost-effectiveness estimates for cash transfers are available at www.givewell.org/international/technical/criteria/cost-effectiveness/costeffectiveness-models

confident in the impact of GiveDirectly's work than in that of any of the other charities discussed in this post.

Transparency and communication.

GiveDirectly has always communicated clearly and openly with us. It has tended to raise problems to us before we ask about them, and we generally believe that we have a very clear view of its operations. We feel more confident about our ability to keep track of future challenges than with any of the other charities discussed in this post.

Risks

GiveDirectly has scaled (and hopes to continue to scale)
quickly. Thus far, it has significantly increased the
amount of money it can move with limited issues as
a result. The case of staff fraud (discussed in our full
review of the organization) that GiveDirectly detected
is one example of an issue possibly caused by its pace of
scaling, but its response demonstrated the transparency
we expect.

Schistosomiasis Control Initiative (SCI)⁸

Important changes in the last 12 months

As discussed above regarding DtWI, Dr. Kevin Croke released a new study of a randomized controlled trial of a deworming

⁸GiveWell's full review of SCI is available at http://www.givewell.org/international/top-charities/schistosomiasis-control-initiative

program showing large, long-term impacts from deworming programs. This study is a significant positive update on the impacts of deworming and increased our confidence that deworming programs have significant long-term impacts.

We continued our work revisiting SCI's case for impact. There appear to have been major problems with some, though not all, of the studies we had relied on (pre-2013) to assess SCI's impact. SCI shared some additional monitoring information with us which supported the conclusion that its programs have generally succeeded, though these reports have significant limitations.

We also reviewed the papers of several academics who had previously been critical of SCI's activities. We found little in this literature to change our views on SCI's programs.

We spent significantly more time with SCI in 2014 (including a 3-day visit to its headquarters in London) than we had in previous years, aiming to improve our understanding of its operations and spending. The picture that emerged was more detailed though largely consistent with what we believed before. Specifically:

We are less confident in our understanding of how SCI has spent unrestricted funds. At the end of 2013, we believed we had a relatively strong understanding of SCI's unrestricted spending, but after spending additional time reviewing reports and discussing with SCI staff, we have more questions than we did a year ago.

 We have better information about how SCI plans to use additional funds it receives and the constraints, besides funding, that SCI faces in utilizing additional funding.

Funding gap

SCI told us that it has approximately \$3.8 million worth of opportunities that it would be highly likely to undertake if it had the funding available. (Some of this would be spent in 2015 and some held for the following year to ensure programs can continue once started). It believes it could possibly absorb an additional \$4.5 million (up to \$8.3 million total) for opportunities that are more speculative. Overall, our best guess is that SCI will use up to approximately \$6.3 million and, beyond that, would build up reserves.

Partly for reasons of donor coordination, we have set its target at \$6.8 million total. We hope that SCI will receive \$1 million from individual donors (excluding Good Ventures) this giving season as a result of our recommendation.

Key considerations

Program impact and cost-effectiveness.

Our best guess is that deworming is roughly as cost-effective as distributing bednets and more cost-effective than cash transfers by a factor of 5-10.9

⁹GiveWell's cost-effectiveness estimates for deworming are available at www.givewell.org/international/technical/criteria/cost-effectiveness/costeffectiveness-models

Directness and robustness of the case for impact.

We have seen some evidence demonstrating that SCI successfully deworms children, though this evidence is relatively thin. Nevertheless, deworming is a relatively straightforward program, and we think it is likely (though far from certain) that SCI is successfully deworming people. We have had difficulties communicating with SCI, which has reduced our ability to understand it; we have also spent significant time interviewing SCI staff and reviewing documents over the past 5 years and have found minor but not major concerns.

Transparency and communication.

We have had consistent difficulties communicating with SCI. Specifically, (a) we had a major miscommunication with SCI about the meaning of its self-evaluations ¹⁰ and (b) although we have spent significant time with SCI, we remain unsure of how SCI has spent funds and how much funding it has available (and we believe SCI itself does not have a clear understanding of this). Importantly, if there is a future unanticipated problem with SCI's programs, we don't feel confident that we will become aware of it; this contrasts with AMF and GiveDirectly, both of which we feel we have a strong ability to follow up.

¹⁰We discuss SCI's external communication of its evidence of impact at blog.givewell.org/2014/10/26/update-on-scis-evidence-of-impact.

Risks

There are significantly more unknown risks with SCI than our other top charities due to our limited understanding of its activities. We hope for SCI to have \$6.8 million available, which is significantly more unrestricted funding than it has had available in the past.

Summary

The table below summarizes the key considerations for our four top charities.

Consideration	AMF	DtWI	GiveDirectly	SCI
Estimated	5-10x	10-20x	1x	5–10x (and
cost-effectiveness				possibly
relative to cash				more)
transfers				
Directness and	Strong	Weakest	Strongest	Moderate
robustness of the				
case for impact				
Transparency and	Strong	Strong	Strongest	Weakest
communication				
Ongoing	Strong	Strong	Strongest	Weakest
monitoring				
and likelihood of				
detecting further				
problems				
Organization's	Moderate	Moderate	Strong	Strong
track record of				
rolling out				
program				
Room for more	High	Limited	Very high	Limited
funding				

Table 9.1: Relative strengths and weaknesses of GiveWell's Top Charities

Standouts

Much of the work we did this year went into investigating potential new additions to our top charities list. The strongest contenders we found are discussed below.

Ultimately, none of these made it into our top tier of recommendations, but that could easily change in the future. We believe that more investigative effort could result in a much better understanding of GAIN-USI and potentially a top-tier recommendation. Meanwhile, ICCIDD and DMI do not have the track record we'd want to see for our top tier of recommendations, but in both cases we expect major developments in the next year. Specifically, ICCIDD will have a substantially larger working budget (due to GiveWell money moved), and DMI may have new data from its randomized controlled trial that could cause a significant upgrade in its status.

These are all strong giving opportunities, and we've vetted them all relatively thoroughly. Two work on a program (universal salt iodization) that we believe has excellent cost-effectiveness and a strong evidence base, and the other two have recently released data from randomized evaluations of their own programs. We have thoroughly vetted each of these organizations, including site visits. And we can see arguments for supporting these organizations in lieu of our top charities this year, though we ultimately recommend our top charities above them.

Below are some brief comments on each standout organization. Donors interested in learning more should read our full reviews of each organization.

Development Media International

DMI produces radio and television broadcasts in developing countries that encourage people to adopt improved health practices, such as exclusive breastfeeding of infants and seeking treatment for symptoms associated with fatal diseases. Its programs reach many people for relatively little money, so if its program successfully changes listeners' behavior, it may be extremely cost-effective. It is in the midst of running a randomized controlled trial of its program; the midline results were released earlier this year.

At midline, the study found moderate increases (relative to the control group) in self-reported health behaviors. Our attempt to estimate the likely mortality impact of these behaviors — when accounting for other concerns about the generalizability of the study — implied cost-effectiveness worse than AMF's. This isn't sufficient for a recommendation this year, as DMI has much less of a track record than our top charities. However, if endline results hit DMI's targeted mortality impact, we would expect to adjust our estimate significantly, and DMI could become a top charity.

DMI's current budget is approximately \$2.5 million; it has told us it expects to receive approximately \$2.5–\$4 million from existing funders in the next year and could absorb an

additional \$6-\$7.5 million, which it would either use to supplement a program already broadcasting in a country or move into a new country, depending on how much it received.

GAIN-USI

GAIN's Universal Salt Iodization (USI) program supports salt iodization programs. There is strong evidence that salt iodization programs have a significant, positive effect on children's cognitive development, and we consider the program to accomplish (very roughly speaking) comparable good per dollar to bednets and deworming.

GAIN-USI does not work directly to iodize salt; rather, it supports governments and private companies to do so, which could lead to leveraged impact of donations or to diminished impact depending on its effectiveness. We tried but were unable to document a demonstrable track record of impact; we believe it may have had significant impacts, but we are unable to be confident in this with what we know now. More investigation next year could change this picture.

GAIN's USI program was one of the recipients of a large, multi-year grant from the Bill and Melinda Gates Foundation. The grant ends in 2015 and has yet to be renewed; we are unsure of whether it will be.

Donors whose primary interest is supporting a strong intervention, and who are comfortable supporting a large and reputable organization whose role is to promote and support the intervention (but whose track record we cannot assess at this time), should strongly consider supporting GAIN's USI program.

GAIN is a large organization running many programs, so donors should consider the possibility that funds restricted to GAIN's USI program might effectively support its other efforts. GAIN told us that it has very little unrestricted funding, so it is unlikely to be able to reallocate funds from other programs to continue to support USI work. It is possible that resources that are shared across programs (such as some staff) could be shifted toward other programs if resources for USI increased, but we would guess that this effect would be small.

International Council for the Control of Iodine Deficiency Disorders Global Network.

Like GAIN-USI, ICCIDD supports (via advocacy and technical assistance rather than implementation) salt iodization, and as with GAIN-USI, we tried but were unable to establish a track record of successfully contributing to iodization programs. Unlike GAIN-USI, ICCIDD is small, operating on a budget of approximately half a million dollars per year, and relies heavily on volunteer time. We believe that additional funding in the range of a few hundred thousand dollars could have a significant positive impact on its operations.

Good Ventures has granted a total of \$350,000 to ICCIDD this year, and we would be happy to see ICCIDD receive a few hundred thousand dollars more, after which point we

would be more hesitant as it would be more than doubling its budget. We hope that ICCIDD will use the additional funding to improve its capacity and potentially become a top charity in the future.

Living Goods

Living Goods recruits, trains, and manages a network of community health promoters who sell health and household goods door-to-door in Uganda and Kenya and provide basic health counseling. They sell products such as treatments for malaria and diarrhea, fortified foods, water filters, bed nets, clean cook stoves and solar lights.

It completed a randomized controlled trial of its program and measured a 27% reduction in child mortality. We estimate that Living Goods saves a life for roughly each \$10,000 it spends, approximately 3 times as much as our estimate for the cost per life saved of AMF's program. Living Goods has been operating on a budget of \$3 million per year and aims to scale up to operate on a budget of \$10 million per year, of which it expects to receive approximately two-thirds from existing funders.

Our research process in 2014

This year, we completed an investigation of one new intervention, salt iodization. We made substantial progress on several others (maternal and neonatal tetanus immunization

campaigns, mass drug administration for lymphatic filariasis, and vitamin A supplementation) but did not complete them.

We also stayed up to date on the research for bednets, cash transfers and deworming and made a substantial update to our view on deworming, based on a new study by Kevin Croke.

We did not conduct an extensive search for new charities this year. We feel that we have a relatively good understanding of the existing charities that could potentially meet our criteria, based on past searches. Instead, we solicited applications from organizations that we viewed as contenders for recommendations. (Living Goods is an exception; it contacted us with the results from its randomized controlled trial.)

In addition to the 4 standout charities, we also considered Nothing but Nets (a bednets organization that declined to participate in our process), Evidence Action's Dispensers for Safe Water program (which is forthcoming), the Center for Neglected Tropical Disease and UNICEF's maternal and neonatal tetanus program. In the case of the latter two, we ran out of time to complete the relevant intervention reports this year (due to prioritizing other work, which seemed more likely to lead to new recommendations) and plan to complete them in 2015.

Part III CAREER CHOICE

10

To Save the World, Don't Get a Job at a Charity: Go Work on Wall Street

William MacAskill 1

Few people think of finance as an ethical career choice. Top undergraduates who want to "make a difference" are encouraged to forgo the allure of Wall Street and work in the charity sector. And many people in finance have a mid-career ethical crisis and switch to something fulfilling.

The intentions may be good, but is it really the best way to make a difference? I used to think so, but while researching ethical career choice, I concluded that in fact it's better to earn a lot of money and donate a good chunk of it to the most cost-effective charities-a path that I call "earning to give." Bill Gates, Warren Buffett and the others who have taken the 50% Giving Pledge are the best-known examples. But you don't have to

¹This article was posted by William MacAskill to Quartz on February 27 2013, to Quartz at http://qz.com/57254/to-save-the-world-dont-get-a-job-at-a-charity-go-work-on-wall-street.

be a billionaire. By making as much money as we can and donating to the best causes, we can each save hundreds of lives.

There are three considerations behind this. First is the discrepancy in earnings between the different career paths. Annual salaries in banking or investment start at \$80,000 and grow to over \$500,000 if you do well. A lifetime salary of over \$10 million is typical. Careers in non-profits start at about \$40,000, and don't typically exceed \$100,000, even for executive directors. Over a lifetime, a typical salary is only about \$2.5 million. By entering finance and donating 50% of your lifetime earnings, you could pay for two non-profit workers in your place-while still living on double what you would have if you'd chosen that route.

The second consideration is that "making a difference" requires doing something that wouldn't have happened anyway. Suppose you come across a woman who's had a heart attack. Luckily, someone trained in CPR is keeping her alive until the ambulance arrives. But you also know CPR. Should you push this other person out of the way and take over? The answer is obviously "no". You wouldn't be a hero; you wouldn't have made a difference.

So it goes in the charity sector. The competition for notfor-profit jobs is fierce, and if someone else takes the job instead of you, it's likely that he or she won't be much worse at it than you would have been. So the difference you make by taking the job is only the difference between the good you do, and the good that the other person would have done. The competition for finance jobs is even fiercer than for non-profits, but if someone else gets the finance job instead of you, he or she would be unlikely to donate as much to charity. The average donation from an American household is less than 5% of income — a proportion that decreases the richer the household. So if you are determined to give a large share of your earnings to charity, the difference you make by taking that job is much greater.

The third and most important consideration is that charities vary tremendously in the amount of good they do with the money they receive. For example, it costs about \$40,000 to train and provide a guide dog for one person, but it costs less than \$25 to cure one person of sight-destroying trachoma in the developing world. For the cost of improving the life of one person with blindness, you can cure 1,600 people of it.

This matters because if you decide to work in the charity sector, you're rather limited. You can only change jobs so many times, and it's unlikely that you will be able to work for only the very best charities. In contrast, if you earn to give, you can donate anywhere, preferably to the most cost-effective charities, and change your donations as often as you like.

Not many people consider "earning to give" as a career path. But it's proving popular. We estimate that there are at least 100 people pursuing earning to give in the effective altruism community. One student, convinced by these arguments, now works at Jane Street, the trading firm, giving 50% of his

income, and thus can already pay the wages of several people for the not-for-profit work he could have been doing.

In general, the charitable sector is people-rich but moneypoor. Adding another person to the labor pool just isn't as valuable as providing more money. You might feel less directly involved because you haven't dedicated every hour of your day to charity, but you'll have made a much bigger difference.

11

High Impact Science

Carl Shulman 1

Paul Ehrlich began his 1968 book, The Population Bomb, with this statement:

The battle to feed all of humanity is over. In the 1970s hundreds of millions of people will starve to death in spite of any crash programs embarked upon now. At this late date nothing can prevent a substantial increase in the world death rate.

Ehrlich predicted these deaths as a consequence of the challenge of feeding a rapidly growing world population, pointing to recent devastating famines in South Asia. But even as those words were written, the fields were being planted with new, higher-yielding semi-dwarf strains of wheat and rice. Combined with modern fertilizers and other methods, these strains ushered in the "Green Revolution": wheat production in India and Pakistan almost doubled between 1965 and 1970, and formerly famine-wracked countries became self-sufficient in

 $^{^1{\}rm This}$ article was initially posted to the 80,000 Hours blog in November 2011 at www.80000hours.org/blog/5-high-impact-science.

food and have not seen such hunger since. The agronomist Norman Borlaug, who developed new and more effective methods of plant breeding, used them to develop the key strain of wheat, and brought about expansion of his methods to other crops and deployment in South Asia, played a pivotal role. Some credit him with saving a billion lives, referring to the number of people fed by the increased agricultural production of the Green Revolution.² Indisputably, scientific and technological advances and innovations have brought about almost unbelievable amounts of good.

Not whether, but when?

However, when we ask how we can effectively do good in our careers, the key question is not how much scientific research on the whole has made the world better: it is "how do I expect the world to be different if I take up this career, rather than another one?" If Norman Borlaug had never lived, his discoveries would eventually have been made by others. Continued food scarcity would have evoked both market and government responses in increased research. Fertilizers and other agricultural technologies would have been applied without enhanced crop varieties, capturing some of the Green Revolution's benefits. We should think of achievements like

²Easterbrook, Gregg. "Forgotten benefactor of humanity." ATLANTIC-BOSTON- 279 (1997): 74-82. www.theatlantic.com/magazine/archive/1997/01/forgotten-benefactor-ofhumanity/306101

Borlaug's as bringing about technologies faster, rather than making them possible at all.

However, tremendous impacts are possible through speeding the pace of progress, even slightly. According to the WHO, malaria killed over 781,000 people in 2009.³ If current trends continue, advances in vaccines, bednets, mosquito control and increased deployment efforts will likely eventually drive fatalities down to zero. But leaping ahead in this process by a single year could save 781,000 lives. A single day's speedup would save 2,139 lives. Advancing the process by even 40 seconds would save a life. The question therefore becomes: by how many seconds can you expect to advance your field over your career?

Norman Borlaug advanced progress on a massive issue by years, a truly exceptional achievement that very likely saved millions of lives or more over time. And there are plenty of other scientists, working on important but neglected problems, who could each hope to have an impact comparable to moving malaria eradication forward several days. Such an achievement could do as much good as donating millions of pounds to apply existing vaccines or drugs.

³World Health Organization. "World Malaria Report 2010, Summary. Geneva: WHO; 2010." (2010). www.who.int/malaria/world_malaria_report_2010/malaria2010_summary _keypoints_en.pdf

How much can I advance my field?

When we consider donating to administer treatments for tuberculosis or malaria, we can measure lives saved through randomized experiment: administer the treatment to half of a study population at random, and see how many more members of the vaccinated group remain healthy. However, estimating the impact of individual research contributions is harder.

It's easy to measure inputs such as R&D funding and scientific workforce. For instance, spending on malaria R&D was \$612 million in 2009, double that of 2004 thanks to increased spending by the Gates Foundation and United States NIAID.⁴ So it's fairly easy to find out how much adding your brain or donations to the field will boost the relevant inputs: whether by 10% or by 0.01%. We can then rephrase questions about the impact of additional researchers in terms of the speed increase which results. Will doubling R&D budgets cut development time by at least 10%? 2%? 0.2%? Phrased in this way, it is easier to draw on existing data to judge the plausibility of potential impacts.

In 2011, the RethinkHIV initiative convened a panel to perform cost-benefit analyses of a number of HIV/AIDS interventions. When it turned to HIV vaccine research, it consulted expert HIV researchers to estimate the likelihood that vaccine

⁴PATH Malaria Vaccine Initiative. "Staying the course." Malaria research and development in a time of economic uncertainty (2011). www.malariavaccine.org/files/RD-report-June2011.pdf

development would both be successful and not be pre-empted by substitutes. The cost of vaccine development was estimated based on the costs and frequency of past partially successful vaccine candidates.

They then found very large benefits relative to costs from increasing annual R&D spending from \$900 million to \$1 billion. This increase would accelerate vaccine development by around 5 months (compared to a baseline timeframe of 20-30 years), and as a result vaccine research would outrank the deployment of existing treatments in terms of cost-effectiveness.⁵

Reasoning like this, taking into account the magnitude of potential gains in a field, the scale of inputs so far, track records in similar fields, expert judgment (with caveats) and interim successes, can make a good start at deciding which research areas to bet on when trying to steer towards high impact.

Opportunities for high-impact science are present within a wide variety of fields, including not only public health but also physics, mathematics, computer science, economics, and many more.

⁵Hecht, Robert, et al. "Vaccine Research and Development." (2011). www.copenhagenconsensus.com/sites/default/files/hecht_et_al.pdf

12

How to Assess the Impact of a Career

Ben Todd 1

What does it mean to "make a difference"?

The standard view is that to make a difference with your career, you need to work in organizations with a social mission, such as charities, hospitals and schools. But this view is overly simplistic. Many of the highest-impact people in history were politicians, entrepreneurs, and academics (think Bill Gates or Norman Borlaug). And there's a lot of variation in the good done by different charities: some aren't effective at all.

In order to work out how to make a difference, we need to know what that even means. We propose the following definition: You "make a difference" whenever you contribute to solving important social problems, enabling others to flourish now and in the long-run.

¹This article combines two articles that were posted to 80,000 Hours in 2014: https://80000hours.org/career-guide/framework/ and https://80000hours.org/2014/02/in-which-career-can-you-make-the-most-difference.

Now that we have a working definition, we can begin to compare jobs and careers in terms of the impact they can have. Here we present a framework you can use to compare how much difference you'll make in different career options. Using this framework can help you choose between existing options, and even to find new options you hadn't considered before.

The four factors

You may be choosing between two jobs like working at Oxfam and the Red Cross, or two career paths like consulting and software engineering. Here are the four factors we consider when comparing these options:

1. Role Impact:

Some roles provide better immediate opportunities to solve social problems than others. For example, the President of the United States has a far greater ability to solve major social problems than a used-car salesman.

Role Impact is the extent to which a role will allow you to make a difference immediately, either through giving you influence at a greater scale or enabling you to work on more pressing problems. We typically consider three types of influence: the direct impact you achieve through your work itself; the earnings that you can donate to charity or invest in training; and the potential to advocate for important causes through your work.

2. Career Capital:

Even if a role has little immediate impact, it can still help you to take better positions later in your career, enabling you to make a greater difference in the future. Even the President had to start somewhere!

Career Capital is the extent to which a role enables you to amass valuable skills, connections and credentials to get into better positions in the long run. These resources are especially important when they can transfer to a wide range of careers, allowing flexibility in your future decisions.

3. Personal Fit:

The role itself is important but people also fit certain roles better than others. If you hate children, you won't make a good teacher. However, if you're good at and enjoy what you're looking to do, you're likely to have a greater impact and build better career capital.

Personal Fit is the extent to which you will excel in a role, enabling you to rise to the top.

4. Exploration Value:

Finally, each step you take in your career can teach you more about the options available, enabling you to make better decisions in the future. If you've only worked in consulting, then trying charity work could be a useful way to test out where to spend the rest of your career.

Exploration Value is the extent to which an option will help you to learn about your options so you can make better decisions in the future in order to maximize your impact.

All of the factors are important throughout your career, but the emphasis changes over time. Especially at the beginning, the value of exploration is particularly high. The focus can then shift to building up your career capital, and finally to applying that capital to solving the most pressing problems.

So in which careers can you make the most difference?

Which paths offer the best personal fit and exploration value depends on your unique situation. However, we've started to look for paths that seem especially promising for role impact and career capital. Of those we've investigated so far, some especially promising, though often competitive, career paths to bear in mind are:

Goal	Careers to consider		
Building career	Entrepreneurship, consulting, economics PhD		
capital			
Earning to give	High-end finance, tech entrepreneurship,		
	law, consulting, medicine, engineering,		
	actuarial work, dentistry		
Direct impact	Party politics, founding effective non-profits		
	working inside governments, foundations or		
	international organizations e.g. UN, (some)		
	academic research		
Advocacy	Party politics, journalism, working in foundations,		
	international organizations or policy-oriented		
	civil service		

Table 12.1: Some types of high impact careers

Our research is constantly being updated. See our most up-to-date version at https://80000hours.org/careerguide/profiles/top-careers/.

Part IV CAUSE SELECTION

13

Your Dollar Goes Further Overseas

GiveWell 1

We understand the sentiment that "charity starts at home", and we used to agree with it, until we learned just how much less effective US charity is than charity aimed at the poorest people in the world.

Helping deprived people in the US usually involves tackling extremely complex, poorly understood problems. Many popular approaches simply don't work, and many others have never been investigated, beyond the stories and anecdotes.

In the poorest parts of the world, people suffer from very different problems. A child may die of malaria for lack of a \$10 bednet, or of diarrhea for lack of a 5 cent packet of nutrients.

¹This article was initially posted on GiveWell's Giving 101 series from 2011 at www.givewell.org/giving101/Your-dollar-goes-further-overseas.

The table below illustrates the difference, comparing USfocused charities to international charities.

Type of	Organization	Cost	Impact
Intervention			
Developing	Against	Approximately	Improves health,
-world	Malaria	\$3,300	saves lives
health	Foundation	per life saved	
Early	Nurse	\$10,000 per	Increases academic
childhood	-Family	child served	performance and
care and	Partnership		reduces criminal
education			behavior
US Education	KIPP	\$7,500-\$17,000	Improves academic
		per student per	performance
		year (including	
		state funds)	
Employment	The HOPE	\$10,000 per	Unclear, if any
Assistance	Program	client served	
(NYC)			

Table 13.1: The difference in impact between US-focused and international charities

We estimate that it costs the Against Malaria Foundation approximately \$3,300 to save a human life. This includes transportation, administration, etc. Compare that with even the best US programs: the Nurse-Family Partnership and KIPP both cost over \$10,000 per child served, and their impact is encouraging but not overwhelming.

This is not to say that developing-world aid is easy or simple. Some activities are highly proven and cost-effective; others have very poor track records. As in the US, generating evidence of impact (not just stories) is essential.

14

The Haste Consideration

Matt Wage ¹

People who dedicate a large part of their life to strategically doing as much good as possible — i.e. effective altruists — are able to accomplish vastly more good than most people will. Unfortunately, not many people are effective altruists.

One way to try to improve the world would be to try to convince more people to be effective altruists. If you spent all of your efforts doing this, how long do you think it would take to convince one person who will be at least as effective as you are at improving the world? For most people, if they're strategic about it, I think they could do it in less than two years.

Now imagine two worlds:

(1) You don't do anything altruistic for the next two years and then you spend the rest of your life after that improving the world as much as you can.

¹This article was initially posted to the 80,000 Hours blog in April 2012 at www.80000hours.org/blog/43-the-haste-consideration.

(2) You spend the next two years influencing people to become effective altruists and convince one person who will be at least as effective as you are at improving the world. (Let's assume that this person wouldn't have done anything altruistic otherwise.) You do nothing altruistic after the next two years, but the person you convinced does at least as much good as you did in (1).

By stipulation, world (2) is improved at least as much as world (1) is because, in (2), the person you convinced does at least as much good as you did in (1).

Many people object to this. They think, "It's possible that world (1) could be improved more than world (2) is. For example, world (1) would be better if, in that world, you convinced 10 people to be effective altruists who are at least as good as you." This is a natural thought, but remember that we are assuming that the person you convince in (2) is "at least as good as you are at improving the world". This implies that if you convince 10 people in world (1), then the person you convinced in world (2) will do something at least as good as that. It's true by definition that world (2) is improved at least as much as world (1) is.

There are two lessons we can take away from this. The first lesson is that influencing people to become effective altruists is a pretty high value strategy for improving the world. For any altruistic activity you're doing, it might be useful to ask yourself, "Do I really think this will improve the world more than influencing would?"

The second lesson is that you can do more good with time in the present than you can with time in the future. If you spend the next two years doing something at least as good as influencing people to become effective altruists, then these two years will plausibly be more valuable than all of the rest of your life. In particular, these two years will be more valuable than any two-year period in the future. This is one way to see that the haste consideration is true.

One implication of the haste consideration: It's plausible that how you spend the next few years of your life is more important than how you spend your life after that. For this reason, when choosing a career, you should pay special attention to how each career would require you to spend the next few years. For example, if a career would require you to spend the next few years studying in school and doing nothing altruistic, then this is a major cost of that career.

15

Preventing Human Extinction

Nick Beckstead, Peter Singer & Matt Wage ¹

Many scientists believe that a large asteroid impact caused the extinction of the dinosaurs. Could humans face the same fate?

It's a possibility.

NASA has tracked most of the large nearby asteroids and many of the smaller asteroids. If a large asteroid were found to be on a collision course with Earth, that could give us time to deflect the asteroid. NASA has analyzed multiple options for deflecting an asteroid in this kind of scenario, including using a nuclear strike to knock the asteroid off course, and it seems that some of these strategies would be likely to work.

The search is, however, not yet complete. The new B612 foundation has recently begun a project to track the remaining asteroids in order to "protect the future of civilization on this

 $^{^1{\}rm This}$ article was initially posted to the Effective Altruism blog in August 2013 at www.effective-altruism.com/preventing-human-extinction.

planet." Finding one of these asteroids could be the key to preventing a global catastrophe.

Fortunately, the odds of an extinction-sized asteroid hitting the earth this century are low, on the order of one in a million. Unfortunately, asteroids aren't the only threats to humanity's survival. Other potential threats stem from bioengineered diseases, nuclear war, extreme climate change, and dangerous future technologies.

Given that there is some risk of humanity going extinct over the next couple of centuries, the next question is whether we can do anything about it. We will first explain what we can do about it, and then ask the deeper ethical question: how bad would human extinction be?

The first point to make here is that if the risks of human extinction turn out to be "small," this shouldn't lull us into complacency. No sane person would say, "Well, the risk of a nuclear meltdown at this reactor is only 1 in 1000, so we're not going to worry about it." When there is some risk of a truly catastrophic outcome and we can reduce or eliminate that risk at an acceptable cost, we should do so.

In general, we can measure how bad a particular risk is by multiplying the probability of the bad outcome by how bad the outcome would be. Since human extinction would, as we shall shortly argue, be extremely bad, reducing the risk of human extinction by even a very small amount would be very good.

Humanity has already done some things that reduce the risk of premature extinction. We've made it through the cold war and scaled back our reserves of nuclear weapons. We've tracked most of the large asteroids near Earth. We've built underground bunkers for "continuity of government" purposes, which might help humanity survive certain catastrophes. We've instituted disease surveillance programs that track the spread of diseases, so that the world could respond more quickly in the event of a large-scale pandemic. We've identified climate change as a potential risk and developed some plans for responding, even if the actual response so far has been lamentably inadequate. We've also built institutions that reduce the risk of extinction in subtler ways, such as decreasing the risk of war or improving the government's ability to respond to a catastrophe.

One reason to think that it is possible to further reduce the risk of human extinction is that all these things we've done could probably be improved. We could track more asteroids, build better bunkers, improve our disease surveillance programs, reduce our greenhouse gas emissions, encourage non-proliferation of nuclear weapons, and strengthen world institutions in ways that would probably further decrease the risk of human extinction. There is still a substantial challenge in identifying specific worthy projects to support, but it is likely that such projects exist.

So far, surprisingly little work has been put into systematically understanding the risks of human extinction and how

best to reduce them. There have been a few books and papers on the topic of low-probability, high-stakes catastrophes, but there has been very little investigation into the most effective methods of reducing these risks: we know of no in-depth, systematic analysis of the different potential strategies. It follows that a reasonable first step toward reducing the risk of human extinction would be to investigate these issues more thoroughly, or support others in doing so.

If what we've said is correct, then there is some risk of human extinction and we probably have the ability to reduce this risk. There are a lot of important related questions, which are hard to answer: How high a priority should we place on reducing the risk of human extinction? How much should we be prepared to spend on doing so? Where does this fit among the many other things that we can and should be doing, like helping the global poor? Does the goal of reducing the risk of extinction conflict with ordinary humanitarian goals, or is the best way of reducing the risk of extinction simply to improve the lives of people alive today and empower them to solve the problem themselves?

We won't try to address those questions here. Instead, we'll focus on this question: how bad would human extinction be?

One very bad thing about human extinction would be that billions of people would likely die painful deaths. But in our view, this is, by far, not the worst thing about human extinction. The worst thing about human extinction is that there would be no future generations.

We believe that future generations matter just as much as our generation does. Since there could be so many generations in our future, the value of all those generations together greatly exceeds the value of the current generation.

Considering a historical example helps to illustrate this point. About 70,000 years ago, there was a supervolcanic eruption known as the Toba eruption. Many scientists believe that this eruption caused a "volcanic winter" which brought our ancestors close to extinction. Suppose that this is true. Now imagine that the Toba eruption had eradicated humans from the earth. How bad would that have been? Some 3000 generations and 100 billion lives later, it is plausible to say that the death and suffering caused by the Toba eruption would have been trivial in comparison with the loss of all the human lives that have been lived from then to now, and everything humanity has achieved since that time.

Similarly, if humanity goes extinct now, the worst aspect of this would be the opportunity cost. Civilization began only a few thousand years ago. Yet Earth could remain habitable for another billion years. And if it is possible to colonize space, our species may survive much longer than that.

Some people would reject this way of assessing the value of future generations. They may claim that bringing new people into existence cannot be a benefit, regardless of what kind of life these people have. On this view, the value of avoiding human extinction is restricted to people alive today and people who are already going to exist, and who may want to have children or grandchildren.

Why would someone believe this? One reason might be that if people never exist, then it can't be bad for them that they don't exist. Since they don't exist, there's no "them" for it to be bad for, so causing people to exist cannot benefit them.

We disagree. We think that causing people to exist can benefit them. To see why, first notice that causing people to exist can be bad for those people. For example, suppose a woman knows that if she conceives a child during the next few months, the child will suffer from multiple painful diseases and die very young. It would obviously be bad for her child if she decided to conceive during the next few months. In general, it seems that if a child's life would be brief and miserable, existence is bad for that child.

If you agree that bringing someone into existence can be bad for that person and if you also accept the argument that bringing someone into existence can't be good for that person, then this leads to a strange conclusion: being born could harm you but it couldn't help you. If that is right, then it appears that it would be wrong to have children, because there is always a risk that they will be harmed, and no compensating benefit to outweigh the risk of harm.

Pessimists like the nineteenth-century German philosopher Arthur Schopenhauer, or the contemporary South African philosopher David Benatar accept this conclusion. But if parents have a reasonable expectation that their children will have happy and fulfilling lives, and having children would not be harmful to others, then it is not bad to have children. More generally, if our descendants have a reasonable chance of having happy and fulfilling lives, it is good for us to ensure that our descendants exist, rather than not. Therefore we think that bringing future generations into existence can be a good thing.

The extinction of our species — and quite possibly, depending on the cause of the extinction, of all life — would be the end of the extraordinary story of evolution that has already led to (moderately) intelligent life, and which has given us the potential to make much greater progress still. We have made great progress, both moral and intellectual, over the last couple of centuries, and there is every reason to hope that, if we survive, this progress will continue and accelerate. If we fail to prevent our extinction, we will have blown the opportunity to create something truly wonderful: an astronomically large number of generations of human beings living rich and fulfilling lives, and reaching heights of knowledge and civilization that are beyond the limits of our imagination.

Speciesism

Peter Singer ¹

When we say that all human beings, whatever their race, creed, or sex, are equal, what is it that we are asserting? Those who wish to defend hierarchical, inegalitarian societies have often pointed out that by whatever test we choose it simply is not true that all humans are equal. Like it or not we must face the fact that humans come in different shapes and sizes; they come with different moral capacities, different intellectual abilities, different amounts of benevolent feeling and sensitivity to the needs of others, different abilities to communicate effectively, and different capacities to experience pleasure and pain. In short, if the demand for equality were based on the actual equality of all human beings, we would have to stop demanding equality. Still, one might cling to the view that the demand for equality among human beings is based on the actual equality of the different races and sexes.

 $^{^{\}rm 1}{\rm This}$ chapter is excerpted from the first chapter of Peter Singer's Animal Liberation, the 2002 edition.

However, there is no need to pin the case for equality to one particular outcome of a scientific investigation. The appropriate response to those who claim to have found evidence of genetically based differences in ability among the races or between the sexes is not to stick to the belief that the genetic explanation must be wrong, whatever evidence to the contrary may turn up; instead we should make it quite clear that the claim to equality does not depend on intelligence, moral capacity, physical strength, or similar matters of fact. Equality is a moral idea, not an assertion of fact. There is no logically compelling reason for assuming that a factual difference in ability between two people justifies any difference in the amount of consideration we give to their needs and interests. The principle of the equality of human beings is not a description of an alleged actual equality among humans: it is a prescription of how we should treat human beings.

Jeremy Bentham, the founder of the reforming utilitarian school of moral philosophy, incorporated the essential basis of moral equality into his system of ethics by means of the formula: "Each to count for one and none for more than one." In other words, the interests of every being affected by an action are to be taken into account and given the same weight as the like interests of any other being.

It is an implication of this principle of equality that our concern for others and our readiness to consider their interests ought not to depend on what they are like or on what abilities they may possess. Precisely what our concern or consideration requires us to do may vary according to the characteristics of those affected by what we do: concern for the well-being of children growing up in America would require that we teach them to read; concern for the well-being of pigs may require no more than that we leave them with other pigs in a place where there is adequate food and room to run freely. But the basic element — the taking into account of the interests of the being, whatever those interests may be — must, according to the principle of equality, be extended to all beings, black or white, masculine or feminine, human or nonhuman.

Thomas Jefferson, who was responsible for writing the principle of the equality of men into the American Declaration of Independence, saw this point. It led him to oppose slavery even though he was unable to free himself fully from his slaveholding background. He wrote in a letter to the author of a book that emphasized the notable intellectual achievements of Negroes in order to refute the then common view that they had limited intellectual capacities:

Be assured that no person living wishes more sincerely than I do, to see a complete refutation of the doubts I myself have entertained and expressed on the grade of understanding allotted to them by nature, and to find that they are on a par with ourselves...but whatever be their degree of talent it is no measure of their rights. Because Sir Isaac Newton was superior to others in understanding, he was not therefore lord of the property or persons of others.

It is on this basis that the case against racism and the case against sexism must both ultimately rest; and it is in accordance with this principle that the attitude that we may call "speciesism," by analogy with racism, must also be condemned. If possessing a higher degree of intelligence does not entitle one human to use another for his or her own ends, how can it entitle humans to exploit nonhumans for the same purpose?

It may be objected that comparisons of the sufferings of different species are impossible to make and that for this reason when the interests of animals and humans clash the principle of equality gives no guidance. It is probably true that comparisons of suffering between members of different species cannot be made precisely, but precision is not essential. Even if we were to prevent the infliction of suffering on animals only when it is quite certain that the interests of humans will not be affected to anything like the extent that animals are affected, we would be forced to make radical changes in our treatment of animals that would involve our diet, the farming methods we use, experimental procedures in many fields of science, our approach to wildlife and to hunting, trapping and the wearing of furs, and areas of entertainment like circuses, rodeos, and zoos. As a result, a vast amount of suffering would be avoided.

Four Focus Areas of Effective Altruism

Luke Muehlhauser 1

It was a pleasure to see all major strands of the effective altruism movement gathered in one place at the recent Effective Altruism Summit.

Representatives from GiveWell, The Life You Can Save, 80,000 Hours, Giving What We Can, Animal Charity Evaluators, Leverage Research, the Center for Applied Rationality, and the Machine Intelligence Research Institute either attended or gave presentations. My thanks to Leverage Research for organizing and hosting the event!

What do all these groups have in common? As Peter Singer said in his TED talk, effective altruism "combines both the heart and the head". The heart motivates us to be empathic and altruistic toward others, while the head can "make sure that

¹This article was previously posted to the rationality blog LessWrong in July 2013, and refers to the Effective Altruism Summit from earlier that year (www.lesswrong.com/lw/hx4/four_focus_areas_of_effective_altruism/).

what [we] do is effective and well-directed," so that altruists can do not just some good but as much good as possible.

Below are four popular focus areas of effective altruism, ordered roughly by how large and visible they appear to be at the moment. Many effective altruists work on several of these focus areas at once, due to uncertainty about both facts and values.

Though labels and categories have their dangers, they can also enable chunking, which has benefits for memory, learning, and communication. There are many other ways we might categorize the efforts of today's effective altruists; this is only one categorization.

Focus area 1: Poverty reduction

Here, "poverty reduction" is meant in a broad sense that includes (e.g.) economic benefit, better health, and better education.

Major organizations in this focus area include the following:

- GiveWell is home to the most rigorous research on charitable causes, especially poverty reduction and global health.
- Good Ventures works closely with GiveWell.

- The Life You Can Save (TLYCS), named after Peter Singer's book on effective altruism, encourages people to pledge a fraction of their income to effective charities. TLYCS currently recommends GiveWell's recommended charities and several others.
- Giving What We Can (GWWC) undertakes some charity evaluation and also encourages people to pledge 10% of their income to effective charities. GWWC currently recommends two of GiveWell's recommended charities and two others.
- AidGrade evaluates the cost effectiveness of poverty reduction causes, with less of a focus on individual organizations.

In addition, some well-endowed foundations seem to have "one foot" in effective poverty reduction. For example, the Bill & Melinda Gates Foundation has funded many of the most cost-effective causes in the developing world (e.g. vaccinations), although it also funds less cost-effective-seeming interventions in the developed world.

In the future, poverty reduction effective altruists might also focus on economic, political, or research-infrastructure changes that might achieve poverty reduction, global health, and educational improvements more indirectly, such as when Chinese economic reforms lifted hundreds of millions out of poverty. Though it is generally easier to evaluate the cost-effectiveness of direct efforts than that of indirect efforts,

some groups (e.g. the Open Philanthropy Project and The Vannevar Group) are beginning to evaluate the likely cost-effectiveness of these causes.

Focus area 2: Meta effective altruism

Meta effective altruists focus less on specific causes and more on "meta" activities such as raising awareness of the importance of evidence-based altruism, helping effective altruists reach their potential, and undertaking research to help them decide where to focus their efforts.

Organizations in this focus area include the following:

- 80,000 Hours highlights the importance of helping the world effectively through your career. They also offer personal counseling to help effective altruists choose a career and a set of causes to support.
- Explicitly, the Center for Applied Rationality (CFAR) just trains people in rationality skills. But de facto they are especially focused on the application of rational thought to the practice of altruism, and are deeply embedded in the effective altruism community.
- Leverage Research focuses on growing and empowering the effective altruism movement, e.g. by running the Effective Altruism Summit, by organizing the THINK student group network, and by searching for "mind

hacks" (like the memory palace) that can make effective altruists more effective.

Other people and organizations contribute to meta effective altruism too. Paul Christiano examines effective altruism from a high level at his Rational Altruist blog. GiveWell and others often write about the ethics and epistemology of effective altruism in addition to focusing on their chosen causes. And, of course, most effective altruist organizations spend some resources growing the movement.

Focus area 3: The far future

Many effective altruists value future people roughly as much as currently-living people, and therefore think that nearly all potential value is found in the well-being of the astronomical numbers of people who could populate the far future. Future-focused effective altruists aim to somewhat-directly capture these "astronomical benefits" of the far future, e.g. via explicit efforts to reduce existential risk.

Organizations in this focus area include:

The Future of Humanity Institute at Oxford University is the primary hub of research on existential risk mitigation within the effective altruism movement. Recently, The Centre for Study of Existential Risk in Cambridge and the Future of Life Institute in Boston have joined the existential risk research effort...)

The Machine Intelligence Research Institute focuses on undertaking the research necessary for humanity to one day build Friendly Artificial Intelligences that could make astronomical numbers of future people enormously better off. It also runs the LessWrong group blog and forum, where much of today's effective altruist analysis and discussion occurs.

Other groups study particular existential risks (among other things), though perhaps not explicitly from the view of effective altruism. For example, NASA has spent time identifying nearby asteroids that could be an existential threat, and many organizations (e.g. GCRI) study worst-case scenarios for climate change or nuclear warfare that might result in human extinction but are more likely to result in "merely catastrophic" damage.

Some effective altruists (e.g. Holden Karnofsky, Paul Christiano) have argued that even if nearly all value lies in the far future, focusing on nearer-term goals (e.g. effective poverty reduction or meta effective altruism) may be more likely to realize that value than more direct efforts.

Focus area 4: Animal suffering

Effective animal altruists are focused on reducing animal suffering in cost-effective ways. After all, animals vastly outnumber humans, and growing numbers of scientists believe that many animals consciously experience pleasure and suffering. The only organization of this type so far (that I know of) is Animal Charity Evaluators, which currently recommends supporting Animal Equality International, Mercy for Animals and The Humane League.

Major inspirations for people in this focus area include Peter Singer, David Pearce, and Brian Tomasik.

Other focus areas

I could perhaps have listed "effective environmental altruism" as focus area 5. The environmental movement in general is large and well-known, but I'm not aware of many effective altruists who take environmentalism to be the most important cause for them to work on, after closely investigating the above focus areas. In contrast, the groups and people named above tend to have influenced each other, and have considered all these focus areas explicitly. For this reason, I've left "effective environmental altruism" off the list, though perhaps a popular focus on effective environmental altruism could arise in the future.

Other focus areas could later come to prominence, too.

Working together

I was pleased to see people from different strands of the effective altruism movement cooperating and learning from each other at the Effective Altruism Summit. Cooperation is crucial for growing the movement, so I hope that even if it's

Four Focus Areas of Effective Altruism

not always easy, effective altruists will "go out of their way" to cooperate and work together, no matter which focus areas have their sympathies.

Part V ORGANIZATIONS

GiveWell

GiveWell

What Is GiveWell?

GiveWell is a non-profit dedicated to finding outstanding giving opportunities to help donors decide where to give. We publish the full details of our analysis at our website, www.givewell.org.

How is GiveWell different from other "charity evaluators" (e.g., Charity Navigator)?

Unlike charity evaluators that focus solely on financials, assessing administrative or fundraising costs, we conduct indepth research aiming to determine how much good a given program accomplishes (in terms of lives saved, lives improved, etc.) per dollar spent. Rather than try to rate as many charities as possible, we focus on the few charities that stand out most by our criteria in order to find and confidently recommend the best giving opportunities possible.

What are GiveWell's criteria for evaluating giving opportunities?

We use different criteria for the two different parts of our research: our traditional work on finding evidence-backed international aid charities and our newer project "The Open Philanthropy Project".

Evidence-backed International Aid Charities

Our current top charities are characterized by the following qualities:

- Serving the global poor. Low-income people in the developing world have dramatically lower standards of living than low-income people in the US, and we believe that a given dollar amount can provide more meaningful benefits when targeting the former.
- Focused on evidence-backed interventions. We have a high standard for evidence: we seek out programs that have been studied rigorously and repeatedly, and whose benefits we can reasonably expect to generalize to large populations (though there are limits to the generalizability of any study results). The set of programs fitting this description is relatively limited, and mostly found in the category of health interventions (though there is also substantial evidence on cash transfers).

• Thoroughly vetted and highly transparent. We examine potential top charities thoroughly and skeptically, and publish thorough reviews discussing both the strengths of these charities and any concerns. We also follow our top charities' progress over time and report on it publicly, including any negative developments. Charities must be open to our intensive investigation process — and public discussion of their track record and progress, both the good and the bad — in order to earn "top charity" status. (We also provide a list of charities meeting our first two criteria for donors who are concerned that this requirement creates problematic selection effects.)

Open Philanthropy Project

More recently, GiveWell has been increasing the breadth of our research as part of the Open Philanthropy Project. In this project, we're open-among other things-to funding political advocacy, scientific research, startup organizations with no track record, projects with no precedent, and projects with extremely long time horizons.

We do not yet have any giving recommendations in these areas, and the project is evolving rapidly. The rest of this document focuses on our current top charities; more information on the Open Philanthropy Project is available at http://www.open-philanthropy.org.

What Is GiveWell's Research Process?

Thoroughly investigating even a small number of charities requires a great deal of work; thousands of hours of research have gone into our recommendations. Our research process generally includes reviewing the independent research behind charities' programs, researching possible concerns about these programs, conducting extensive back-and-forth with charities to fully understand their processes and past and future uses of funds, time-intensive cost-effectiveness analysis (estimating how much good is accomplished per dollar spent), multi-day site visits to charities' operations in the field, and publishing ongoing updates on charities' activities.

What Are GiveWell's Current Top Charity Recommendations?

GiveWell's current top charities (as of March 2014) are:

- Against Malaria Foundation
- GiveDirectly
- Schistosomiasis Control Initiative (SCI)
- Deworm the World Initiative (led by Evidence Action)

Our list of top charities is refreshed annually.

Why Should I Trust GiveWell's Recommendations?

Donors can thoroughly evaluate GiveWell's recommendations for themselves because GiveWell is committed to extreme transparency: we publish the full details of the reasoning behind our recommendations, notes from phone conversations, audio recordings of our Board meetings, a blog with regular content on our evolution, and more. We believe that information about how to help people should never be kept secret.

What is GiveWell's impact?

In 2013, GiveWell tracked \$17.36 million in donations to our top charities as a direct result of our research. Our research has received accolades from Peter Singer (who featured it in his manifesto on giving, *The Life You Can Save*), Nicholas Kristof (who featured it in *Half The Sky*), the Hewlett Foundation (a current funder of GiveWell), and media including the *New York Times, Wall Street Journal*, CNN and NPR.

How Can I Learn More?

Explore www.givewell.org, follow our blog and sign up for our email list.

Giving What We Can

Michelle Hutchinson, Executive Director

Introduction

Giving What We Can is an international community whose members pledge to give a significant fraction of their income to whichever charitable organizations can do the most good. We also conduct research aimed at identifying those organizations that are most effective in improving the lives of people in the world's poorest countries.

The Pledge to Give

Each of our members takes a pledge to give at least 10% of their income until retirement. By giving to the very best charities, we can make extraordinary improvements to the lives of poor people in developing countries. Recognising that we can each live well enough with less, our members make a commitment to give to whichever organisations can most effectively improve the welfare of others, now and in the years to come. For those who don't yet feel ready to commit to giving 10%, we

also run a *Try Giving* scheme. This allows you to commit to donating a percentage of your choosing over a short period, thereby easing into the idea of donating a significant amount over the long-term.

An International Community

We created an international community because we can achieve far more if we stand together and make a public pledge. Signing a pledge acts as a form of pre-commitment, an effective psychological tool for following through on important life-goals. As a community, we can support each other in many different ways, such as by sharing valuable information about charity cost-effectiveness. And in making our giving public, we can inspire others to join the cause. Giving What We Can has garnered significant media attention from the very beginning, appearing in the New York Times, the Guardian, the Times, the Wall Street Journal, to name but a few.

Research

We've always emphasized the importance of evaluating charities in terms of what really counts: helping people to the greatest extent possible. Evidence suggests that most of the variance in the effectiveness of charities is due to the kind of intervention they implement: whether they are building schools, treating trachoma, providing microloans, and so on. For this

reason, we apply a top-down research procedure that begins by examining different cause-areas and ends with careful assessments of charities that are exemplary in carrying out the best of programmes. As of the start of 2015, Giving What We Can recommends Against Malaria Foundation, Deworm the World, Project Healthy Children, and Schistosomiasis Control Initiative.

Impact

We have already achieved an incredible amount. Giving What We Can was founded in 2009 by Toby Ord and Will MacAskill. Our membership has since grown to over 900 people. We have established chapters and local meetup groups in the UK, Switzerland, Germany, the US, and Australia, with more many more in the pipeline.

By joining Giving What We Can, you can become part of this dynamic, growing community of effective altruists. Most importantly, by giving 10% of your income to the very best charities, you can make an extraordinary difference to countless lives that would otherwise be blighted by malaria, malnutrition, and other symptoms of extreme poverty. Between us, we've already donated nearly \$7 million to highly effective charities. The total amount pledged by our members is currently \$386 million and it's growing fast. We can't wait to see what the future holds.

The Life You Can Save

Charlie Bresler, Executive Director

Following the 2009 publication of his book, The Life You Can Save, Peter Singer, along with some of his colleagues and philosophy students at Oxford, Princeton, and Rutgers, started an organization by the same name. The goal of The Life You Can Save (TLYCS) was to obtain annually-renewed pledges from people to donate a percentage of their annual income (the percentage increasing the more they earn) to charities that are effectively engaged in fighting the devastating effects of extreme poverty among the 1.2 billion poorest people in the world — those living on less than \$1.25 USD/day. Since that time almost 17,000 people have taken the pledge.

Currently, TLYCS has two full-time associates, two parttime associates, a volunteer Executive Director, several other key volunteers, and an annual operating budget of approximately \$200,000 USD. The current goal is to move a minimum of five times the operating budget to the NGOs recommended on TLYCS website. The longer-term goal is to mass market the ideas embedded in Peter's work, stressing the ethical obligation to share your good fortune with the least fortunate, to new audiences, thus increasing the pool of donors contributing to the fight against extreme poverty.

TLYCS continues to encourage people to pledge an annual amount and runs giving games for university students (2,000 students thus far) to introduce them to effective altruism: participants decide, from a pre-determined list of effective charities, where to donate their money. We expect that the lifetime value of a percentage of these participants' donations will grow significantly every year, as these students graduate, enter the workforce, donate, pledge, and, renew their pledges annually. TLYCS also supports "pledge groups that contact us, or that we start. We leverage the marketing materials we create and curate, as well as our relationships with researchers investigating how to maximize commitment to fighting extreme poverty, across these pledge groups.

TLYCS uses a concept of "personal best" to augment the pledge strategy. This strategy, which is widely supported in the social psychology literature, as well as in competitive athletics, assumes that if individuals who are not yet ready to take the pledge are encouraged to donate more than they have previously donated, and to more effective NGOs, the overall amount of money donated effectively will grow more than if the pledge strategy alone is employed. The idea is that using "personal best" will dramatically grow the pool of individuals who join the effective altruist community, since it does not

require as large a commitment as a lifetime pledge, or even an annual pledge.

The goal of mass marketing and the use of "personal best" differentiates TLYCS from other groups in the effective altruism community and augments the highly effective work of groups like Giving What We Can, which has raised a large amount of money for effective NGOs through a 10% lifetime pledge and an intensive, rather than an extensive, focus like TLYCS. We hope that all these groups, working cooperatively, will help in the fight to eliminate extreme poverty.

80,000 Hours

Ben Todd, Co-founder and Executive Director

Introduction

Although over a third of students prioritize making a difference with their careers, advice on how to do so is poor. Because of this, every year over half a million graduates fail to make as much difference with their careers as they could.

Our proposed solution is to become the best source of advice in the world for these talented graduates, so they can make far more of a difference with their careers. Through this, create a global, evidence-based, systematic conversation about how to best use your career to solve the world's most pressing problems.

Our vision is for as many people as possible take high impact careers.

What we do

Research and online content

We carry out in-depth research with academics at the University of Oxford into the best strategies and opportunities for making a difference with your career. We publish this on our research pages.

Coaching

Our research is driven by one-on-one coaching. We currently receive about 30 requests for coaching every month. We aim to speak to around a third of these, and undertake one indepth case study. We let the real choices facing graduates guide the issues we explore.

Community

By making introductions through coaching and maintaining an online members' directory, we're building a community of people who want to make the greatest possible difference through their careers.

Our plans

Over the next year we will complete our discovery phase, in which our aim is to identify the most high potential business model. Our focus is on improving the prototype of our online content and the depth of the research behind it. We're doing this by writing answers to the most pressing questions for the people we coach, and preparing career profiles on the most promising careers. Based on what we learn, we'll re-write the research pages and submit them for evaluation.

In August 2015, Will MacAskill's book on Effective Altruism will be released. Guardian Faber in the UK has already pledged \$60,000 towards advertising. This will start our outreach phase, during which we'll aim to reach as much of our audience as possible through the media and events.

The effective altruist case for 80,000 Hours

Only a few years ago, donors had almost no useful information about the impact of different charitable projects. Today, young altruists are similarly in the dark about the best strategies for maximizing the impact of their careers.

By providing this information, we can enable many more people to find and take the most high impact career opportunities. In this way, we'll act as a *flexible resource multiplier* on the best causes, with far more impact than we could have had directly.

We think that the flexibility of this strategy is key. We're uncertain about which interventions will be most high impact in the future; but 80,000 Hours will be able to direct people to them, whatever they turn out to be.

In terms of our impact so far, we believe we've significantly changed the career plans of over 100 talented young people. We think this already more than justifies historical financial costs of under \$200,000.

If this project interests you...

Then you may like to:

- Read our research at 80000hours.org/career-guide
- Apply for coaching at 80000hours.org/career-advice

Charity Science

Xiomara Kikauka, Co-founder

What is our goal?

At Charity Science we are taking all of the amazing evidence already out there and making it accessible to a wider audience than just the people who get excited by spreadsheets. In essence, we are the popular science of charity work.

What activities do we do?

Charity Science mainly focuses on outreach and educating the public about the science of doing good. To do this we:

- Run fundraising experiments to raise money for the most scientifically proven charities.
- Give presentations to the public, companies and schools.
- Build a community by organizing dinners, pub nights, potlucks, and meet ups where like-minded people can get to know each another.

- Run events to raise awareness and funds for the most evidence-based, cost effective, and transparent charities.
- Write articles about research into charities and how to make an effective charity.

We organize our activities into "experiments", not "projects". We generate hypotheses about which activities will help the world and then test them. We then rigorously measure and check the data to evaluate the activities.

Most charities run "projects" and when a "project" doesn't work out, it often feels like a failure. This biases philanthropists to try to find a rationalization to continue the "project", so they don't feel like they've failed. This causes many charities to continue to work on ineffective activities.

This isn't the case with experiments. If the result of the experiment is that an activity does not help the world, then you've learned something important about the world, even if it means working no further in the area.

Why scientifically proven charities?

Science has shown again and again that intuition and reasoning are often wrong. Psychological research indicates that we cannot predict what will make ourselves happy; what makes us think that, sitting in our armchairs, we can know what will make somebody on the other side of our city happy? Or on the other side of the world?

Charity Science

The world is complicated. We cannot simply assume that things that "make sense" will work in the real world. We need to get out of our armchairs, test our theories and then use statistics to see what works.

How can I get involved?

We would love to get volunteers to help with outreach, research, and in many other areas. You can be local (in Vancouver, Canada) or international.

For more information on running an effective charity and helping the world generally, or to get involved with the Charity Science, visit our website at www.charityscience.com

23

Effective Altruism Foundation

Adriano Mannino, President and Jonas Vollmer, Executive Director

The Effective Altruism Foundation (EAF, "EA-Stiftung" or "EAS" in German) is an EA non-profit and think tank based in Switzerland. It acts as the umbrella organisation for several projects launched by effective altruists in Switzerland, Germany and Austria. More recently, EAF also set up international projects in collaboration with effective altruists from overseas. Its activities have a common focus on "metacharity": growing the movement and researching the best strategies for doing good, with an emphasis on reducing suffering. Until mid-2015, EAF was known as "GBS Switzerland".

Raising for Effective Giving

EAF has been able to become a major influence in the international poker scene through its project Raising for Effective Giving (REG). REG was launched in collabora-

tion with famous poker professionals such as Liv Boeree and promotes Earning to Give and other EA ideas in the poker community. EA outreach targeted at communities and industries that already share a rational, expected-value-oriented mindset can be particularly effective. In 2014, REG fundraised more than half a million dollars for effective charities (poverty/animals/AI/meta), with total expenses of about \$50,000. REG was thus able to function as a ten-fold multiplier meta-charity in 2014. More information can be found on our website: www.reg-charity.org

Sentience Politics

With the project Sentience Politics ("Politics for all sentient beings"), EAF tries to create sizeable impacts in the areas of direct animal charity, animal meta-charity and exploration of grassroots and institutional-political strategies. One demand we've launched in the Swiss public sphere, using various political instruments, is the introduction of more vegetarian and vegan options in public cafeterias. We've been able to gain the support of a former minister and several members of parliament. Among other things, submitting controversial but interesting political demands can draw a lot of societal attention to the organisation behind them – Sentience Politics got month-long coverage and was featured in a documentary on national TV. Among other additional things, we're promoting the ethical importance of donating and Earning to Give in the animal movement, and are spreading awareness about

neglected high-stakes areas such as wild animal suffering, insect suffering and potential (animal-like, "voiceless") artificial suffering. For more information: www.sentience-politics.org

Effective Altruism Outreach

In addition to outreach efforts on its blogs and on social media, EAF hosts public talks and runs chapters at several universities. Through these channels we are trying to spread EA ideas to a larger, promising audience as well as to present avenues for getting involved. EAF organises regular meet-ups and provides a local network for people interested in optimising their professional careers for altruistic value.

Philosophy and Research

EAF's focus is on reducing the amount of involuntary suffering in the world in peaceful, sustainable ways. This includes contemporary human suffering, the suffering of non-human animals both on factory farms and in the wild, the suffering of future generations as well as potential future suffering resulting from artificial intelligence and artificial consciousness. Prioritising among these areas (and all the possible interventions) is a hard task, but respective work is likely to be very valuable if progress can be made.

For this reason, EAF set up the Foundational Research Institute (FRI). FRI tries to identify the best ways to do good, with a focus on preventing dystopian scenarios for the world. Recently FRI has highlighted compromise between different

value systems, international cooperation and the prevention of technological arms races as potential priorities for future study. FRI informs the strategic decision-making at EAF, translating academic insights into actual policies and outreach efforts. For instance, EAF published a position paper on artificial intelligence in mid-2015, co-authored with neurophilosopher Prof. Thomas Metzinger.

Further Information

Visit the website www.ea-stiftung.org for more information on the projects and activities, and www.foundational-research.org for an overview of the philosophical assumptions and strategic research that motivate and inform EAF's altruistic efforts.

24

The Machine Intelligence Research Institute

Luke Muehlhauser, Executive Director

The Machine Intelligence Research Institute (MIRI) was founded in 2000 on the premise that creating smarter-than-human Artificial Intelligence with a positive impact — "Friendly AI" — might be a particularly efficient way to do as much good as possible.

First, because future people vastly outnumber presently existing people, we think that:

"From a global perspective, what matters most (in expectation) is that we do what is best (in expectation) for the general trajectory along which our descendants develop over the coming millions, billions, and trillions of years."

¹Nick Beckstead details the case for this claim in his thesis On the Overwhelming Importance of Shaping the Far Future. The same kind of is made in this book's chapter on Preventing Human Extinction.

Second, as an empirical matter, we think that smarter-than-human AI is humanity's most significant point of leverage on that "general trajectory along which our descendants develop." If we handle advanced AI wisely, it could produce tremendous goods which endure for billions of years. If we handle advanced AI poorly, it could render humanity extinct. No other future development has more upside or downside.²

Third, we think that Friendly AI research is tractable, urgent, and uncrowded.

Tractable: Our staff researchers and visiting workshop participants tackle open problems in Friendly AI theory, such as: How can we get an AI to preserve its original goals even as it learns new things and modifies its own code? How do we load desirable goals into a self-modifying AI? How do we ensure that advanced AIs will cooperate with each other and with modified versions of themselves? This work is currently at a theoretical stage, but we are making clear conceptual progress, and growing a new community of researchers devoted to solving these problems.

Urgent: Surveys of AI scientists, as well as our own estimates, expect the invention of smarter-than-human AI in the second half of the 21st century if not sooner. Unfortunately, mathematical challenges such as those we need to solve to build Friendly AI often require several decades of research to overcome, with each new result building on the advances

² As argued in Nick Bostrom's Superintelligence: Paths, Dangers, Strategies.

that came before. Moreover, because the invention of smarterthan-human AI is so difficult to predict, it may arrive with surprising swiftness, leaving us with little time to prepare.

Uncrowded: Very few researchers, perhaps fewer than five worldwide, are explicitly devoted to full-time Friendly AI research.

The overwhelming power of machine superintelligence will reshape our world, dominating other causal factors. Our intended altruistic effects on the vast majority of beings who will ever live must largely reach them via the technical design of the first self-improving smarter-than-human Als. Many ongoing efforts — on behalf of better altruism, better reasoning, better global coordination, etc. — will play a role in this story, but we think it is crucial to also directly address the core challenge: the design of stably self-improving Als with desirable goals. Failing to solve that problem will render humanity's other efforts moot.

If our mission appeals to you, you can either fund our research at www.intelligence.org/donate or get involved in other ways at www.intelligence.org/get-involved.

25

Animal Charity Evaluators

Jon Bockman, Executive Director

Animal Charity Evaluators (ACE) is a non-profit 501(c)(3) charitable organization located in the United States.

What is ACE's goal?

Our mission at ACE is to find and advocate highly effective opportunities for improving the lives of animals. To do this, we analyze the impact of interventions (i.e. tactics) to help animals. Our general definition of an "effective intervention" is one that creates comparatively larger positive impact (measurable reduction of suffering or alternatively an improvement of well-being) for sentient beings at low cost. By extension, we evaluate animal charity organizations based on how well they meet the ultimate goal of effectiveness.

The practical implication of our work is twofold: one, we educate individuals (especially activists or altruists) about the most effective ways to make a difference for animals; and two, we recommend the highest-performing charities, based on our findings, to effectiveness-minded donors.

We use research from other organizations and conduct our own research to gather the largest possible pool of data to inform our decisions. Currently, very little reliable research exists around effective animal advocacy. Our aim is to help close that gap directly but also encourage other organizations to produce more high-quality research in order to continuously build a shared body of knowledge and ever-improving findings.

ACE looks to aid compassionate givers, professionals, and volunteers alike in making informed decisions on how to be as effective as possible in helping animals. We strive to identify highly effective opportunities for alleviating suffering and improving animal lives on a wide scale, and will continue to update our recommendations based on new evidence.

Why focus on animals?

We certainly believe there is great value in addressing global poverty and other human-centric causes. Nevertheless, given that our goal is to most efficiently reduce the largest amount of suffering, there is a very compelling case to focus on animals. Animals vastly outnumber humans. To illustrate, almost 60 billion animals are bred and killed for food each year worldwide, compared to the entire human population of 7.13 billion (as of December 2013). Furthermore, as evidenced by

numerous sources (undercover videos, standard agricultural practices), the suffering that animals endure — often caused by human self-interest — is enormous. By educating people about how they can best advocate for animals, we contribute to the largest possible reduction of suffering for the largest number of sentient creatures.

By addressing the current system of institutionalized animal exploitation, particularly the food industry, from the viewpoint of helping animals, we acquire major wins for the human population (health, global hunger, etc.). Additionally, a unique characteristic of farm animal advocacy is that it translates into everyone's everyday food choices, causing significant potential flow-through effects. Not only does farm animal advocacy directly lower the demand for animal products (and in turn lower supply), but also each changed individual will likely perpetuate this change through influencing others around them, sometimes without any additional organizational intervention. Lastly, shifting purchasing power away from factory farm products immediately creates new demand — and in turn, supply — for ethical alternatives, reinforcing the change.

How did you come to your conclusions about top interventions and top charities?

We begin our process by defining the initial framework of interventions and charities based on the following criteria.

We start with interventions which have demonstrated the potential to affect large numbers of animals at low cost, and which either have direct evidence of effectiveness or which allow for relatively easily-performed studies to provide initial evidence for or against their effectiveness. Then we primarily investigate charities which perform one of our recommended interventions or which publicly provide other evidence that their own activities have comparable cost-effectiveness to our highest-performing interventions.

For interventions and charities that pass the initial screening process, we use our methodically designed intervention template and charity recommendation template to guide more thorough investigations. We use existing data, including data provided by the charities, and perform our own studies to estimate effectiveness where necessary. Although we use quantitative and empirical data from authoritative sources wherever possible, our charity recommendations ultimately involve subjective factors, including how to weigh competing criteria. We publish in-depth discussions of our research so that donors can understand our reasoning and substitute their own judgments if their values differ from ours.

Note that we continuously refine and perfect this overall process, learning from our successes and mistakes as well as from newly available external findings. It is possible that new developments will lead to us redefining our criteria in the future. As this process is not an exact science, we never guarantee bullet-proof results, but we believe in extra trans-

parency and we always disclose our methodology, especially the subjective factors.

Why should I trust ACE's recommendations?

ACE offers a unique perspective in the animal advocacy movement by using objective means to identify the social efficacy of animal welfare tactics and charities. We hold no stake in any one group or intervention, and our sole interest in promoting the best ways to help animals ensures that we do not hold a bias toward any specific area.

We update our top charities on December 1st of each year. Outside of this cycle, we occasionally update our recommendations if warranted by significant new research or other reliable information.

Afterword

Ryan Carey

Once you think about it, effective altruism can seem like an obvious idea. In the early days of the effective altruism community, newcomers were often surprised that so many people didn't *already* analyze their positive impact on the world. Over the past few years, the community has grown so rapidly that newcomers are now greeted by many others who share their perspective.

In a few short years, thousands of pledges have been made, thousands of hours of research have been performed, and hundreds of millions of dollars have been moved. Effective altruists have saved thousands of people's lives in the developing world.

Hopefully, effective altruism will continue to grow in popularity.

In particular, I hope that this book can help people to gather up their excitement for this philosophy, to support these valuable organizations and, ultimately, to make a bigger difference in the world.