



**CENTRE FOR
EFFECTIVE
ALTRUISM**

Centre for Effective Altruism

Fundraising Prospectus – Winter 2017

Contents

Executive Summary	3
2016 Retrospective and Plans for 2017	4
Community & Outreach Division	6
Effectivealtruism.org	6
Effective Altruism Global conferences	7
Giving What We Can pledge	9
Giving What We Can Trust	10
Chapters	11
Pareto Fellowship	13
Doing Good Better	13
Individual Outreach	14
Special Projects Division	15
Fundamentals research	16
Oxford Institute for Effective Altruism	17
Philanthropic Advising	18
Policy Research and Consulting	20
Mistakes We've Made	21
Budget Summaries	23
CEA Team	25

Executive Summary

The Centre for Effective Altruism (CEA) helps to grow and maintain the effective altruism movement. Our mission is to

- create a global community of people who have made helping others a core part of their lives, and who use evidence and scientific reasoning to figure out how to do so as effectively as possible; and
- make the advancement of the wellbeing of all a worldwide intellectual project, doing for the pursuit of good what the Scientific Revolution did for the pursuit of truth.

We have two divisions. The community and outreach division focuses on growing, strengthening, and serving the effective altruism community. The special projects division focuses on improving our understanding of how to do the most good, which includes exploring new applications for effective altruism.

2016 was a year of significant change for CEA. We went from being a collection of largely autonomous teams to a single team under one management structure. Although the internal reorganization is complete, the process of integrating the various projects continues. Our overarching goal for 2017 is to build a strong, focused CEA, which we believe is essential to achieving our long-term mission.

Key projects we intend to pursue in 2017 include the following:

- Turn effectivealtruism.org into the landing page for the effective altruism community
- Host three Effective Altruism Global conferences
- Establish a scalable model for facilitating student and local groups
- Launch a multidisciplinary institute at the University of Oxford for the study of effective altruism
- Develop advanced quantitative cause prioritization models

Below we provide a detailed review of our activities over the past year and our plans for next year. In summary, we believe the case for supporting CEA based on marginal cost-effectiveness alone is strong. But we also maintain that that is not how CEA should be evaluated.

We believe that effective altruism has the potential to have a transformative impact on how people think about doing good in the world and that CEA is currently best positioned to help effective altruism realize its potential. We estimate that most of CEA's value in expectation comes from the chance—small as it might be—that it realizes that mission. In other words, funding CEA is a gamble, albeit (because of its marginal cost-effectiveness) a gamble in which the “bad” outcome still looks pretty good.

For 2017, the minimum we're looking to raise is £2.5 million. We believe we could spend much more than that before hitting strongly diminishing returns: we could spend £5.1 million in our growth scenario and £7.3 million in our stretch growth scenario. In both of these latter two scenarios, we would regrant a significant amount of money to smaller projects in the effective altruism community.

If you'd like to donate to CEA, you can do so using the options on our donate page: www.centreforeffectivealtruism.org/donate

If you any further questions, please contact Alison Woodman at alison.woodman@centreforeffectivealtruism.org.

2016 Retrospective and Plans for 2017

We feel that 2016 was a good year for CEA. We accomplished a lot, and although we made some mistakes, we took steps to correct our mistakes and lay the foundation for an even better 2017. Below is a review of the past year and a discussion of our plans for next year, first at an organizational level and then at a project level.

Finances. We will have spent £1,332,298 in 2016, broken down as follows:

- £821,463 for CEA UK; of which £555,810 was spent on staff salaries

- \$640,000 [£512,000] for CEA US; of which \$330,130 [£264,104] was spent on the Effective Altruism Global conference and \$251,567 [£201,254] was spent on staff salaries

Staff. We grew from from 18 to 23 full-time-equivalent staff members.

Key accomplishments. Below are some noteworthy accomplishments from 2016. We describe our impact in more detail in our following discussion of specific projects.

- Gained 945 new Giving What We Can pledges, who pledged £357 million to effective charities
- Hosted Effective Altruism Global, the largest ever effective altruism conference (with over 1000 attendees)
- Hosted eight EAGx conferences on five different continents (with four more scheduled for early 2017)
- Launched effectivealtruism.org, which is already generating 15,000 unique sessions per month
- Authored the forthcoming report “Existential Risks - Diplomacy and Governance” with over 50 collaborators on behalf of the Finnish Ministry for Foreign Affairs
- Launched the effective altruism concepts project, a collection of nearly 200 concepts relevant to effective altruism research.

Because so much giving happens in December, and it takes time to record donations made, our most accurate recent “money moved” figure is from 2015, for which we recorded \$6,939,624.40 in donations from our members, of which we believe about \$2.5 million would not have happened were it not for us.

Reorganization. A significant event from an organizational perspective was CEA’s internal reorganization. CEA was previously a collection of largely autonomous projects, comprising Giving What We Can, Effective Altruism Outreach, Global Priorities Project, CEA Central, and 80,000 Hours. As the effective altruism community grew, these projects grew in an often-uncoordinated way. In July of this year, the board of trustees

decided to merge the projects into a single management structure (with the exception of 80,000 Hours).

CEA is now divided into two divisions: a community and outreach division, and a special projects division. The community and outreach division focuses on growing, strengthening, and serving the effective altruism community. It runs givingwhatwecan.org, effectivealtruism.org, Effective Altruism Global, EAGx and the Pareto Fellowship, and facilitates over 70 local groups around the world. The special projects division focuses on improving our understanding of how to do the most good, which includes exploring new applications for effective altruism. It conducts research into effective philanthropy, policy and cause prioritization, and is doing the groundwork to set up an academic institute focused on effective altruism at the University of Oxford.

Goals for 2017. We discuss project-level goals in more detail below, but our high-level goals are the following:

1. Strengthen the foundation of the effective altruism community. The effective altruism community is reaching a size where we should expect dilution, coordination failures, in-fighting, and stagnation, among other problems. Over the next year, we will invest in strengthening the community's foundation, including by developing community resources on effectivealtruism.org.
2. Identify and facilitate new opportunities. The combination of a growing community, new ideas, and increasing contact with influential individuals creates the potential for new high-value opportunities. We believe we are well-positioned to identify and facilitate such opportunities, even if we ultimately leave the implementation to others. We have been actively exploring new opportunities in policy and philanthropy.
3. Focus. We are currently running too many projects while simultaneously looking for new high-value opportunities. Our goal for next year is to review our ongoing projects and scale up those that look promising and scale down or discontinue those that look less promising. We will also prioritize projects that have a synergistic relationship with other, ongoing projects.

4. Establish a strong team culture. Perhaps most important of all, we want to develop a strong team with a shared vision.

Community & Outreach Division

Effectivaltruism.org

In August, the online infrastructure team (Tara Mac Aulay, Kerry Vaughan, Larissa Hesketh-Rowe, Sam Deere, and Oliver Habryka) launched effectivaltruism.org, which contains a new introduction to effective altruism, a collection of the best writing on effective altruism, community profiles, and an interactive cause prioritization tool, among other things. We intend to invest significantly in this site, with the aim of making it an onboarding tool for people new to effective altruism and a valuable resource for community veterans.

Examples of work

- Building a chapters portal, providing chapter leaders with a portal they can update, which pulls events from Facebook to keep it up to date
- Writing and testing an email drip campaign (set of emails sent out on a defined schedule) to encourage new visitors to get more involved and learn core effective altruism ideas
- Creating an interactive tool to guide donation decisions

Impact review

Effectivaltruism.org gets about 15,000 unique hits per month and is the top Google search result for “effective altruism”. About ten new visitors to effectivaltruism.org take the Giving What We Can pledge each month, and about 10% of visitors sign up to the Effective Altruism Newsletter, which results in them receiving our drip campaign and introduction-to-effective-altruism emails. We currently devote about 12% of community and outreach division staff time to the site. We think that if we invest further in developing content and in marketing the website, we can easily double or triple the amount of new visitors to the website

in the next year and ensure that people new to effective altruism are onboarded to the online and in-person communities.

Plans for next year

We have big plans for effectivealtruism.org over the next year, including the following:

- Launch a chapter portal with resources for local groups (90% confident this will happen)
- Launch a community portal with a layered discussion forum and aggregated effective altruism-relevant content from the web (70%)
- Publish a series of long-form essays that methodically explain the foundations of effective altruism (50%)

Effective Altruism Global conferences

In 2016, the events team (Amy Willey Labenz, Julia Wise, and Roxanne Heston) hosted the largest Effective Altruism Global to date, with more than 1,000 attendees, three workshop rooms, five discussion rooms, and three stages hosting about 80 speakers. We also launched the first Effective Altruism Global Research Meeting and the Effective Altruism Global Leaders Forum.

In addition, we started the inaugural year of the EAGx conference series, with eight events thus far, and four more planned for early 2017. The EAGx conferences this year took place across five continents in eight countries.

Impact review

Overall, we think Effective Altruism Global 2016 was a success. The results of our two post-event surveys reflect the following:

- Attendee satisfaction with the conference overall received a score of 6.1 out of 7 in our closing survey and 4.1 out of 5 in our follow-up survey.
- 14% of the approximately 90 respondents to the follow-up survey said that they changed their life plans significantly as a result of the

conference, 27% said that the conference changed how much they intend to give to effective charities, and 5% reported that they took the Giving What We Can Pledge as a result of the conference.

- 23 attendees signed the Giving What We Can Pledge at the conference or in the weeks after the event.

A more detailed impact review of Effective Altruism Global 2016 is linked to from the online version of this document at www.centreforeffectivaltruism.org/fundraising.

We have less data for this year's EAGx events, but the early returns are promising. Conferences for which we have data averaged a 7.74/10 rate of recommendation, and many conference organizers reported interest in future involvement. While we are unsure about many of the effects of the events, we see 2017 as an opportunity to roll out our improved data collection and analysis system and follow up on the longer-term effects we are hoping to promote.

Plans for next year

The events team has already started planning next year's EAG conferences. After considering the lessons from Effective Global 2016, as well as community feedback and survey results, we are planning to have three main events this year with distinct focuses:

1. Boston / Cambridge, Massachusetts (May). This event will likely focus on more speculative causes and interventions with tracks including topics in science, technology, and policy.
2. San Francisco, California (July or August). This event will likely be primarily community-driven, with the majority of programming coming from members of the effective altruism community in the form of talks, workshops, facilitated discussions, and other specialized events.
3. London / Oxford, United Kingdom (October or November). This event will likely have an academic focus that builds on the 2016 Research Meeting and includes discussions of cause prioritization, ethics, and philosophy.

For EAGx, we are making a number of changes for next year, including: (i) making conference resources publicly accessible on the eaglobal.org website; (ii) offering more flexible funding and application deadlines to conference organizers; (iii) creating streamlined processes for getting community members involved with the events and event attendees involved with the community; and (iv) improving our systems for collecting data and sharing our impact evaluations. Since Roxanne will no longer be directing the series, we will spend the first few months of next year onboarding our new director, accepting the first round of 2017 event applications and making our resources and impact tracking publicly available.

Giving What We Can pledge

The Giving What We Can pledge is a lifetime commitment to give 10% of one's income to effective charities. We view it as a pillar of the effective altruism community. Even after the reorganization, we continue to make Giving What We Can pledges one of our core metrics.

Examples of work

The Giving What We Can pledge is not a distinct project but a route to impact and a way of measuring the impact of several projects. For example, we measure the success of projects like Effective Altruism Global, the newsletter, and effectivealtruism.org in part by reference to the number of pledges they generate. We try to directly drive pledges primarily through online marketing, such as retargeting ads. We also run an annual pledge drive, which we launched this year on November 29 (“Giving Tuesday”).

Impact review

In the last year (December 1, 2015 to November 30, 2016), we gained 945 new Giving What We Can pledges, from a base of 1300 members (73% growth), increasing total pledges by \$357 million, from a base of \$485 million (74% growth).

In 2015 our members moved \$6,939,624.40 to charity in total. Of this we estimate, based on members' self-reports, that 36% would not have happened were it not for the activities of CEA.

Plans for next year

Many of our activities which we expect to lead to pledges are also broadly aimed at increasing awareness of and engagement with effective altruism. For example, we project around \$1,500 per month of spending on Facebook retargeting, if we continue to see similar returns to recent months (it costs around \$2 per newsletter signup using this method), and this expands our audience for promoting effective altruism and the pledge. We anticipate we could spend \$140,000 on running a Doing Good Better book giveaway next giving season, based on the popularity of the campaign so far this year, and contingent upon maintaining a reasonable conversion of book handouts to pledges. We're also considering rewarding grants to chapters explicitly to encourage pledges.

Giving What We Can Trust

The Giving What We Can Trust (primarily operated by Alison Woodman) was set up by employees at CEA in 2013 to collect donations and make grants in support of highly effective global poverty organizations. It's a UK-registered charity enabling tax-efficient giving by UK donors, including in support of organizations and projects based outside of the UK. It also allows donors to set up one donation in support of multiple charities and helps CEA to keep track of the donations we've inspired.

Examples of work

- Automated our bookkeeping and reporting
- Worked with auditors and lawyers to ensure that we have the right systems in place to handle large grants both domestically and overseas
- Introduced automatic receipts and begun allowing donors to opt into communications with the organizations they support
- Added several additional organizations that GiveWell classifies as “promising”, prioritizing those that cannot receive Gift Aid directly from UK donors

Impact review

The amounts donated to the Trust have grown substantially:

- 2014: £440,000 by 708 donors
- 2015: £1.2 million by 1923 donors
- £367,000 came from Giving What We Can Members (£63,000 of which has not been reported by members in My Giving, who joined during or prior to 2015)
- Q1-Q3 2016: £1.3 million by 1765 donors
- £329,000 came from Giving What We Can Members (£96,000 of which has not been reported in My Giving)

We track the channels through which donors reach the online donation forms, and find that donations in 2016 came from the following sources (by percentage of total donations):

- 59% are made by donors visiting www.givingwhatwecan.org directly
- 21% are from GiveWell referrals
- 10% are referrals from other charities' websites
- 10% through payroll giving and other online donation services

Plans for next year

In 2017, we intend to increase our grant-making capacity and expand the number of organizations to which the Trust can donate, including organizations working in areas beyond global poverty.

Chapters

The chapters team (Peter Buckley, Jon Courtney, and Harri Besceli) supports student and local effective altruism groups around the world. Our goal is to grow in-person, local communities and particularly to influence college students who might be willing to choose effective career paths.

Examples of work

We help groups by:

- Skyping regularly with group organizers
- Distributing funding
- Distributing physical resources like books and t-shirts
- Collecting and distributing useful online guides and marketing materials
- Skyping with promising individuals within groups
- Facilitating communication between groups
- Organizing cross-group campaigns like pledge drives
- Connecting individuals within groups to the larger EA network
- Connecting groups with EA-related speakers and workshop providers

Impact review

There are currently over 70 student and local effective altruism groups worldwide, with more starting each month. We believe that most of the value of these groups comes from how their members will use their careers to improve the world, which makes measuring impact challenging. Nonetheless, members of student and local groups have already gone on to work for most major organizations in the effective altruism community, including GiveWell, CEA, 80,000 Hours, Effective Altruism Foundation, and Animal Charity Evaluators. Moreover, this fall, student activities fairs alone led to more than 5,000 new subscribers to the newsletter.

Plans for next year

In 2017, we plan to continue our existing support activities and hope to significantly increase the amount of funding and physical resources we provide to groups. We also hope to increase interaction between groups and group leaders, perhaps by organizing retreats and meetups. If we

increase the size of the team, we will spend more time traveling to chapters and interacting individually with chapter members and leaders.

Pareto Fellowship

The Pareto Fellowship (Tyler Alterman and Peter Buckley) was a new, experimental summer program designed to equip fellows with skills that might increase their prospects for succeeding at a range of world-improving projects. We received nearly 500 applications and narrowed this number down to 18 fellows. The program consisted of two parts: a training period and a projects period.

The training period was four weeks long and focused on topics like effective altruism, movement building, interpersonal skills, model-building, coordination, and epistemic and instrumental rationality. The projects period was four weeks long, during which the fellows co-worked on individual projects while receiving daily one-on-one coaching and support from staff.

A detailed review of the Pareto Fellowship is forthcoming.

Plans for next year

We do not plan to continue the Pareto Fellowship in its current form this year. While we thought that it was a valuable experiment, the cost per participant was too high relative to the magnitude of plan changes made by the fellows. We might consider running a much shorter version of the program, without the project period, in the future. The Pareto Fellowship did, however, make us more excited about doing other high-touch mentoring and training with promising members of the effective altruism community.

Doing Good Better

We sometimes promote *Doing Good Better*, authored by Will MacAskill, as a way to promote effective altruism. In July of this year, the paperback came out in the UK and US. The German edition was published in April this year, and it will be published in South Korea next year.

Examples of work

- Provided copies of *Doing Good Better* for free at Effective Altruism Global in Berkeley and EAGx Oxford.
- Ran a *Doing Good Better* giveaway, which has already generated 2400 book requests and over 2000 nominations of friends who should read the book.
- Will appeared on the Sam Harris podcast, which was downloaded over 400,000 times. He was also featured in *Esquire*, *Spear's*, *Third Sector*, *Die Zeit*, *Tagesspiegel*, *Deutschlandradio Kultur*, *Kulturzeit*, *ARTE TV*, and *Dagens Nyheter*.

Impact review

In the past 12 months, we sold approximately 15,000 copies in the UK, 4,800 in the US, and 3,100 in Germany, not including copies bought directly from the publisher for distribution at Effective Altruism Global events. Since July 2015, 143 new Giving What We Can pledgers (12% of total pledges during that time period) reported that they first heard of Giving What We Can by reading the book.

The book has also been the primary means through which we've received introductions to influential individuals. We believe this is where most of the book's expected value is, and as we discuss in the next section below, we intend to invest more in leveraging these opportunities.

Plans for next year

It is likely (>50%) that we will do more writing on topics in effective altruism that could lay the groundwork for a further book. It is unlikely (<20%) that this would result in another book completed in 2017.

Individual Outreach

In addition to our core work, we have taken opportunities to meet with influential individuals who have expressed interest in effective altruism. These opportunities primarily come through Will's public-facing work (e.g., *Doing Good Better*) but are regularly supported by, or have led to opportunities for, other CEA projects, such as the philanthropic advising team.

Examples of work

- Attended or spoke at a number of salons and conferences
- Met with interested individuals one-on-one
- Provided personal philanthropic advice
- Made introductions to relevant experts and organizations

Impact review

In most cases we aren't able to discuss the impact we've had through individual outreach, though we believe that our engagement here has been highly promising. A few examples that we can mention:

- As a result of an introduction we made, a donation of \$700,000 was made to a charity that the donor believed to be four times more effective than the charity to which they otherwise have donated.
- After several meetings with Will and members of the philanthropic advising team, a Giving Pledge member is considering a sizable (c. \$1,000,000 per year) recurring donation to a GiveWell charity.
- In summer of 2015, Will attended Founders Forum (representing 80,000 Hours), a conference for European entrepreneurs. There, we identified Founders Pledge as an unusually promising new meta-charity. Since then, as a result of interaction with us, they made effective altruism a core component of their mission, hiring three people from the effective altruism community. We raised £345,000 for them from three donors and put them on the radar of the Open Philanthropy Project, who made a grant of \$1 million to them, executed via CEA. Founders Pledge has now raised \$180 million in legally binding pledges from entrepreneurs to give to the charities of their choosing.

Plans for next year

We will likely (80%) significantly increase our investment in this area, proactively reaching out to influential individuals who might be interested in effective altruism.

Special Projects Division

Fundamentals research

The Global Priorities Project had two focus areas: research into fundamental issues like cause prioritization, and policy research and consulting. The former area is now the focus of the fundamentals research team. The fundamentals research team prioritizes projects based on estimated impact along three dimensions: (i) progress at the research frontier, (ii) value to the effective altruism community, and (iii) academic value. Stefan Schubert was the fundamental research team's first and only full-time researcher until late October, when he was joined by Max Dalton. Owen Cotton-Barratt, Pablo Stafforini, Amanda Askell and Ben Garfinkel are part-time researchers.

Examples of work

- With help from the entire Special Projects Division, we created a map of the effective altruism research landscape with nearly 200 concept summaries and links to key resources.
- Stefan is writing an overview of considerations pertaining to future investment into catastrophic risk reduction.
- Owen is working on a multi-year grant from the Future of Life Institute regarding prioritization within AI alignment funding.
- Amanda is writing a series of pieces on low-resilience credences and the value of information relative to direct benefits.
- Ben is writing a paper that explores the phenomena of exponential progress over time (Moore's Law) and exponentially diminishing returns over time (Eroom's Law) with a simple mathematical model.
- Max, Ben, and Owen are developing quantitative cause prioritization models.

Impact review

It's too early to assess the value of this work. Over the next year, we will evaluate research in the same way we prioritize new research projects: by comparing the resources invested to the output generated. How we measure the success of the output will depend on where the projected value is. If the projected value is to the effective altruism community, for example, we will likely consider web traffic and surveys of the targeted part of the community. By contrast, if the projected value is at the research frontier, we will likely ask external experts to assess the value of the work. We expect most of the value to be concentrated in a small percentage of the total research output.

Plans for next year

The priority for next year is to develop a strong internal research culture. CEA has launched multiple research teams over the past five years with mixed results. We are trying to learn from our own mistakes and invest in the foundation of a research team that we can successfully build on in the future.

One component of this goal is looking for ways to narrow our focus without shutting off avenues for potentially creative, high-value work. For example, we will likely deprioritize work focused on influencing academia, particularly when the Oxford Institute for Effective Altruism (discussed below) launches next year.

We will likely continue to do work that is directly useful to the effective altruism community, such as the effective altruism concepts project, but expect to work more closely with the community and outreach division on these projects, and in some cases might play a secondary role.

Example projects in this category include:

- Developing a community research blog (65%)
- Expanding the effective altruism concepts project (50%)
- Creating a database of research projects for students (35%)

We expect to be slow to hire, and will take great caution to hire for the right skills.

Oxford Institute for Effective Altruism

Michelle Hutchinson, with assistance from Jon Courtney, is setting up a multidisciplinary academic institute at Oxford to research topics in effective altruism. Hilary Greaves, professor of philosophy at Oxford University, will lead the institute's research team.

Examples of work

We are primarily focusing on submitting grant applications. To date, we have submitted four grant applications, with one rejection, one probable success, and two on which we've received no information.

Impact review

It is too early to assess the value of the Institute. Ultimately, we hope that it will help increase focus on effective altruism-relevant topics within academia, which we expect would significantly increase the rate of intellectual progress on these topics.

Plans for next year

We are aiming for the Institute to go live in October 2017. Our priorities will be: (i) securing grant funding; (ii) hiring two-to-three postdoctoral researchers, including at least one in economics and one in philosophy; (iii) refining our research agenda; and (iv) launching a website. We are also considering hosting an academic conference (in conjunction with the events team) in Fall 2017.

Philanthropic Advising

Prior to CEA's reorganization, the Giving What We Can research team (Marinella Capriati, Hauke Hillebrandt, and James Snowden) focused part-time on providing bespoke giving advice to wealthy individuals and foundations. Post CEA's reorganization, we've made philanthropic advising a distinct, experimental focus area. The philanthropic advising team is pursuing a two-pronged approach: (i) try to find new promising funding opportunities, and (ii) bring funding opportunities to its network of donors and foundations.

Examples of work

- Produced nine reports for Founders Pledge donors identifying effective charities within specific cause areas
- Advised Alwaleed Philanthropies, a \$30 billion foundation focused on global humanitarianism, regarding focus areas and an evidence-based approach to philanthropy; we ultimately suggested it partner with J-PAL on a large project focused on women's empowerment in the middle east, and put it in touch with J-PAL.
- Advised former Prime Minister Gordon Brown on health education for his International Commission on Financing Global Education Opportunity
- Marinella is investigating promising donation opportunities among organizations that support evidence-based policy in low and middle income countries
- James is investigating pesticide regulation as a promising intervention to reduce suicides
- Hauke is searching for promising funding opportunities in meta research, which aim to improve scientific practice

Impact review

We had mixed results with our pre-reorganization philanthropic advising work. For example, our Founders Pledge clients donated £227,585 to GiveWell-recommended charities (rather than to non-GiveWell-recommended charities), based on our advice. Although we are pleased to have diverted more money to these charities, we had hoped our work would have influenced a larger total amount.

It might be too early to assess the value of some of our client-based work, however. For example, we believe we influenced the Alwaleed Foundation's funding agenda but don't yet know enough to consider the impact of that work. We also think that the Alwaleed Foundation might

(20%) partner with J-PAL on a major new initiative based on our advice, but we have low confidence in this estimate.

We estimate that we had moderate expected impact working with Will on one-on-one outreach to highly influential individuals. As noted above, this work recently led to a Giving Pledge member considering a sizable (c. \$1 million) recurring donation to a GiveWell charity.

It's too early to assess the value of our post-reorganization research into new funding opportunities.

Plans for next year

Whether and in what form we continue our philanthropic advising work will depend on how promising our early research looks. We are likely (75%) to continue exploring new funding opportunities. The primary uncertainties lie in what areas we focus and who the audience is. We will aim to narrow our focus along both dimensions.

We probably will not (35%) focus on providing bespoke research and advice to wealthy individuals and foundations, although we will likely continue this work reactively to take advantage of unusually high-value opportunities.

Policy Research and Consulting

As noted above, Global Priorities Project had two focus areas: research into fundamental issues like cause prioritisation, and policy research and consulting. Since CEA's reorganization, Sebastian Farquhar, in conjunction with Toby Ord, has focused on the latter area. Although we considered various ways to expand our policy focus, we decided against pursuing any of them in the short-term.

Examples of work

- Authored Global Catastrophic Risks 2016 with the Global Challenges Foundation, which was covered on national television, radio, and over 150 newspapers worldwide
- Authored the forthcoming report "Existential Risks - Diplomacy and Governance" with over 50 collaborators on behalf of the Finnish Ministry for Foreign Affairs

- Advised the UK's Department for International Development on their project selection and planning process to increase attention to intervention cost-effectiveness, as well as other areas like pandemic prioritisation
- Advised the Houses of Parliament and White House Office of Science and Technology Policy on the implications of long-term artificial intelligence
- Sebastian and Owen co-authored a forthcoming paper with the Future of Humanity Institute on biosafety pandemic insurance

Impact review

It is difficult to assess policy impact; policy is a collaborative process with many inputs, outcomes emerge over long time-scales, and counterfactuals are hard to assess. Based on conversations with our colleagues in policy roles, we believe our policy work may have had significant impact in the following ways:

- Directly supporting policy-makers to apply cause prioritization in their prioritization process in development, research funding, and risk management affecting billions of annual spending
- Building capacity for catastrophic risk reduction by creating introductory content for government officials, building attention in the international diplomatic community, and supporting the creation of an existential risks training program for international security professionals
- Advising the UK and US governments on the long-term implications of artificial intelligence, after which one of our recommendations (alongside others), a standing body to monitor developments in the area, is likely to be implemented

Plans for next year

We are unlikely (40%) to expand our policy focus over the next year. Although we believe our work in this area has had significant potential impact, we decided not to expand it for the following reasons: (i) it's not clear how we would scale it; (ii) CEA might not be the right organization

to be doing some of this work; and (iii) we believe we can get most of the value of our current work by reactively pursuing high-value opportunities. We would be open to expanding our policy focus, however, if the right potential hire or opportunity presents itself.

Mistakes We've Made

We have of course made mistakes. A few of the more-significant mistakes (and corrections) which we have identified that cut across multiple projects are as follows:

- We were insufficiently transparent. In October, we began publishing monthly detailed updates on our activities on the EA Forum and intend to publish more and more-detailed analyses of various projects.
- In our charity promotion, we paid too little attention to certain cause areas, including animal suffering and the long-run future. In 2017, we intend to promote effective charities that work in these areas and hope to hire staff with expertise in these areas.
- We spread our resources across too many projects. As we note above, focusing our efforts is one of our major goals for 2017.

Below is a non-exhaustive list of additional mistakes, organized by project area:

Effective Altruism Global/EAGx

- Our marketing for Effective Altruism Global was overly aggressive.
- Our communications with participants and speakers were often disorganized. For example, we did not send a save-the-date email sufficiently in advance of the conference.
- We were often too slow to respond to organizers of EAGx conferences.

Chapters

- We were often too slow to respond to leaders of student and local groups.
- We were too slow to distribute funds to student and local groups.
- We had two projects working on student and local groups (Giving What We Can and EA Build), which created confusion.

Community

- We had two projects working on building a community (Giving What We Can and EA Build), which created confusion.
- We did not take any action on several incidents involving disruptive behavior within the community where our involvement likely would have reduced the potential harm.

Research

- Giving What We Can research spent too many resources evaluating the same interventions and organizations that GiveWell was evaluating.
- Global Priorities Project was insufficiently focused, leading to fragmented attention and an unclear medium-term trajectory.


Operations

- Our donation processing system wasn't sufficiently user friendly. For example, we were often too slow to provide donation receipts.
- We occasionally were too slow to respond to general inquiries.

Budget Summaries

These figures compare the baseline budget for 2016 (actual spending Jan-Oct, projected spending Nov-Dec) with the baseline budget for 2017.

Note that these figures represent the minimum amounts required to run CEA at current capacity. Our 'minimum' and 'growth' targets — £2.5m (\$3.1m) and £5.2m (\$6.4m) respectively — are targets for growing CEA beyond current capacity to include more ambitious projects for building and growing the effective altruism community.



Item	2016 (actual)	2017 (projected)
Staff salaries and taxes	757,064	1,029,071
Internship expenses	27,396	38,820
Contractors, design and copywriting	85,173	82,920
Staff equipment and training	21,745	56,760
Travel	61,242	88,600
Marketing and online activities	50,339	162,260
Office rent and expenses	128,675	172,186
Chapter support & grants	22,682	70,220
EAG/EAGx grants	144,584	140,000
Sundries	33,398	19,680
Total (GBP)	£ 1,332,298	£ 1,860,517

CEA Team

Trustees



Will MacAskill
CEO and Trustee



Toby Ord
Trustee



Nick Beckstead
Trustee

Community and Outreach Division



Tara Mac Aulay
Chief Operating Officer,
Community and Outreach
Division



Amy Willey Labenz
Senior Events Manager,
General Counsel



Sam Deere
Full Stack Developer



Kerry Vaughan
Senior Marketing Manager



Alison Woodman
Community Liaison



Larissa Hesketh-Rowe
Marketing and
Communications Specialist



Roxanne Heston
Events Coordinator,
Effective Altruism Global X



Julia Wise
Community Liaison



Peter Buckley
Chapter Outreach
Coordinator

Special Projects Division



Michael Page
Director, Special Projects
Division



Amanda Askeff
Visiting Researcher,
Fundamentals



Marinella Capriati
Researcher, Philanthropic
Advising



Jonathan Courtney
Associate, Oxford Institute fo
Effective Altruism



Owen Cotton-Barratt
Research Advisor,
Fundamentals



Max Dalton
Researcher, Fundamentals



Seb Farquhar
Researcher, Policy and
Philanthropic Advising



Ben Garfinkel
Visiting Researcher,
Fundamentals



Hauke Hillebrandt
Researcher, Philanthropic
Advising



Michelle Hutchinson

Project Manager,
Oxford Institute for
Effective Altruism



Stefan Schubert

Researcher, Fundamentals



James Snowden

Researcher, Philanthropic
Advising



Pablo Stafforini

Researcher, Special Projects