



Global Priorities Project Strategy Overview

The Global Priorities Project helps decision-makers effectively prioritise ways to do good.

We do this both by advising decision-makers on programme evaluation methodology and by developing and encouraging specific policies. We form a bridge between the effective altruism community and major institutions.

February 2015



Executive Summary

Although a very young organisation, the Global Priorities Project has already had a great deal of impact. Our work on unprecedented technological risks has been circulated at the highest levels of government. We contributed to the upcoming update of the UK’s National Risk Assessments for catastrophic events. Our research has been well received in academic circles, and is being prepared for publication.

In 2014, the Global Priorities Project ran as a pilot project within the Centre for Effective Altruism. The successes above were promising enough to merit expansion. With the generous support of our donors, we were able to hire a full-time project manager.

The aim for 2015 is to build on our solid grounding of research, expertise, and relationships to transition towards impact. This will involve a deeper focus on policy, building on the success of our Unprecedented Technological Risks report. In addition, we will be doing more than ever before to make our work relevant for specific audiences and widely distributed.

We look forward to expanding in 2015. We have received dozens of expressions of interest from very strong candidates for research and policy positions. At the moment, funding constraints prevent us from hiring in order to increase the pace and quality of our work. We are planning to raise £100,000 by the end of May to build our reserves and allow us to hire additional staff to improve the quality, utility and quantity of our research.

Sebastian Farquhar
Director

Table of Contents

- Executive Summary2**
- Table of Contents.....2**
- Why prioritisation?3**
- What have we achieved so far?3**
- What will be changing?3**
- What is our project plan?3**
- What is our funding situation?4**
- What is the return on donations to the Global Priorities Project?5**
- What are potential future avenues of research? ...5**
- What are our assets?7**
- The team.....7**

Why prioritisation?

Decision-makers are increasingly good at comparing the effectiveness of programmes focussed on the same metric.¹

However, they often feel unable to weigh up benefits that are very different from each other — like youth programmes against research into cancer treatments, or programmes now and those in the future.² Differences between these types of benefits can be very large, and a better understanding of their relative weight would have a big effect on spending priorities.³ Where we identify these differences, it can suggest specific policy changes.

What have we achieved so far?

Our work so far, roughly 1 person-year of staff time, has had a significant impact on both evaluation methodology and direct policies:

- We published a report on the [risks posed by unprecedented technologies](#). This was widely shared within the highest levels of UK government and became a chapter of the [annual report of the Government Chief Scientific Adviser](#).⁴
- We advised the government during an update of the methodology for the UK's [National Risk Register](#) and [National Security Risk Assessment](#) - which cover catastrophic risks.
- We have had more than a dozen other meetings with senior policy-makers in the UK spanning a wide range of topics.

Our research has been influential with foundations and academics:

- We shared our research on [time-discounting](#) with the global public health community. It was selected to be published as part of a book on the topic.
- We developed a [novel approach](#) to modelling benefits of work in areas of uncertain difficulty which is being used for prioritisation and prepared for publication.
- We engaged with several international organisations including the WHO regarding cost-effectiveness evaluation

What will be changing?

In mid-January, we expanded the core team by hiring a project manager as part of a pivot towards tangible impact from our existing body of research. This will result in some immediate changes, which will prepare us for the possibility of more substantial pivots in the future as we gather better impact data:

- We have identified two focus areas - cause prioritisation and policy, which are detailed below.
- For each workstream we have identified clear measures of success based on impact, reach, and reception and will deprioritise unsuccessful approaches.
- We will build up the Global Priorities Project brand through a [dedicated website](#) and reach out to increase awareness and use of our work.
- We will deprioritise plans to engage further on general issues of Cost-Benefit Analysis (e.g., pure time preference) as this area is increasingly occupied by a large number of competing and well-informed voices.⁵
- We will aim to produce one or two academic papers a year in order to build academic credibility but will not make this a focus

What is our project plan?

Our focus over the coming months will be to establish our reputation as experts in prioritisation. Our work divides into two major streams, each of which tests several approaches for impact. We intend to review impact against each stream after six months in order to decide which activities to continue, and which to stop.

At the moment, the topics we have identified for research vastly exceed our capacity. We look forward to addressing this imbalance.

Cause Prioritisation

When deciding how to allocate their resources, organisations and individuals must prioritise between different cause areas. This is a serious challenge, and many lack the capabilities to handle all the nuances robustly on their own. We will produce decision-relevant guidance on key issues in cause prioritisation.

Timing impact

Foundations and individuals need to decide when to invest and when to act. We have developed a model for

optimal spending over time which we can use to address the issue. For foundations, we will produce guidance on the optimal rate to spend down endowments given the giving opportunities available. For individuals, we will release guidance through 80,000 Hours – a career impact consultancy - regarding timing of career decisions and charitable giving.

Crucial considerations

In some cases, the choices between causes are dominated by key ethical considerations, like time preference. We will produce accessible advice on the impact of these crucial considerations, experimenting with alternative forms of communication such as shareable infographics.

Implied trade-offs

Where organisations make spending decisions, they make implicit trade-offs between diverse benefits. We will make these trade-offs explicit, potentially building off the Copenhagen Consensus, to enable users to contrast the implied trade-offs against their intuitions about what is appropriate. We will offer evidence that people can use to inform their intuitions.

Success criteria

For our cause prioritisation work, success will be measured by the impact we have on how they think about these problems, and ultimately on their decisions. Success criteria:

- Our work is cited by institutions, effective altruist organisations, or individuals
- Our work is read widely
- Our work is judged to be of good quality by people we trust.

Policy

In some cases our expertise will put us in a unique position to recommend specific policies, or improve the methodology used by decision-makers to evaluate policy. We can approach these opportunities in three main ways.

Solicited policy recommendations

Where we feel able to respond to open consultations organised by Government departments, we will offer short comments based on our prioritisation expertise.

Unsolicited policy recommendations

Where we are very confident that we have identified a policy that will significantly raise welfare, we will develop

the idea and encourage organisations with an interest in the specific policy to advocate. We will begin with a paper on mandatory biosafety insurance for biological research labs, which has been well received by biological researchers at conferences.

Unsolicited methodology improvements

Where the opportunity arises, we can offer improvements for evaluation methodology, as we did for National Risk Assessment. We will evaluate the impact of our work in this case, but have no plans to continue this approach in the short term unless a high value opportunity arises, so that we can prioritise activities with a higher marginal value of information.

Success criteria

Our work should be judged by improvement to policy - however, this may be a gradual process. In all cases, we will attempt to only recommend policies which we are confident are actually good policy. In addition, we will judge success not only by the fact that we have traction, but by the fact the consequences are good. This will clearly be difficult to judge and may take many years to become apparent. Success criteria in order of declining importance:

- A change we recommend is implemented
- A change we recommend is implemented with some modifications we dislike
- A change we recommend becomes frequently mentioned as one of the options for a policy area
- We discuss our policy with decision-makers and our ideas are well received
- A change we recommend is covered by the media (ideally favourably)

What is our funding situation?

As of end of January 2015, The Global Priorities Project holds roughly £53,000 in reserves, which maintains a runway until October.

We are targeting 12 months of reserves, which is in line with best practice for donor-funded projects. This would require raising roughly an additional £20,000 to increase reserves in addition to funding on-going expenditure. In order to reach 12 months of reserves at the end of Q2 2015, this would require raising approximately £55,000 in order to maintain the Global Priorities Project at its current size.

Table 1 Budget for 2015 as of December 31st 2014

Line item	Cost in calendar year 2015
Salaries (0.75x researcher, 1x project manager)	42,804
Non-staff expenditure (copy-editing, publications, conferences, travel, etc.)	6,000
Central expenses (office, finance, HR, etc.)	18,387
Contingency	6,719
Total	73,910

However, we would like to hire additional research staff in order to add to our existing research programme and increase the quantity and quality of output. CEA will be hiring in the summer, and so a decision about whether or not to hire additional staff would need to be made at the end of Q2 2015.

The first additional hire would ideally have a background in welfare economics or policy. Given the strength applications in the last hiring round and the interest we have already received, we think it would be fairly easy to hire the right staff given sufficient funding. Hiring an additional researcher at current salaries would add roughly £36,500 per year to costs, including NI and share of central costs. This number would be higher if we have to pay a typical post-doc salary.

We would also like to be able to hire one or two interns this summer. Interns are not paid, but receive accommodation, living costs, and expenses. This would cost the Global Priorities Project approximately £800 per month.

Table 2 - Funding needed by May in order to hire a new researcher

Funding need	Value
Additional funding needed above current reserves currently to end of 2015	16,608
Additional cost of researcher for H2 2015	18,112
Additional cost of intern	2,346
Approximate reserves needed to fund H1 2016	56,160
Total funding needed by May to hire staff	93,225

In order to hire a full time employee to start in July and an intern for two months over the summer, while meeting our objective of 12 months of reserves, we will need to raise approximately £93,000 by the end of Q2 2015.

In order to reach this target, we will be launching a £100,000 fundraising drive to finish in May to coincide with the decision timeline for hiring a new employee. If we are unable to raise the full amount by May we may need to hold off on hiring in this cycle. The effort involved in hiring is high enough that it would be unlikely we could hire in that case before the end of the year.

What is the return on donations to the Global Priorities Project?

Data on backwards looking activities are not yet well developed, and even where available they may not offer a good guide for future returns. A priority for the Global Priorities Project over the coming months is to improve impact tracking.

However, over 2014, with approximately one person-year of activity costing roughly £35,000 (not including volunteered time), we had some [significant achievements](#).

At a very rough level, the return on money spent can be estimated as follows. In the future, however, we might spend money differently, as noted below.

£7,000: A high-production-value policy report on Unprecedented Technological Risks and sharing widely in the British government. Our review panel rated it at 3.6 out of 4 in an aggregate measure of usefulness and quality of analysis and writing. In the future, we will aim to spend more on distribution and publishing after producing such a report.

£700: A research paper or talk, with an average length of 2,500 words. Our review panel gave these at an average aggregate score of 3.1 out of 4. In the future we will spend more on distribution and editing.

£400: A meeting with policy-makers in government including preparation and ground-work, travel and meeting time. This activity is not constrained by money as much as opportunities.

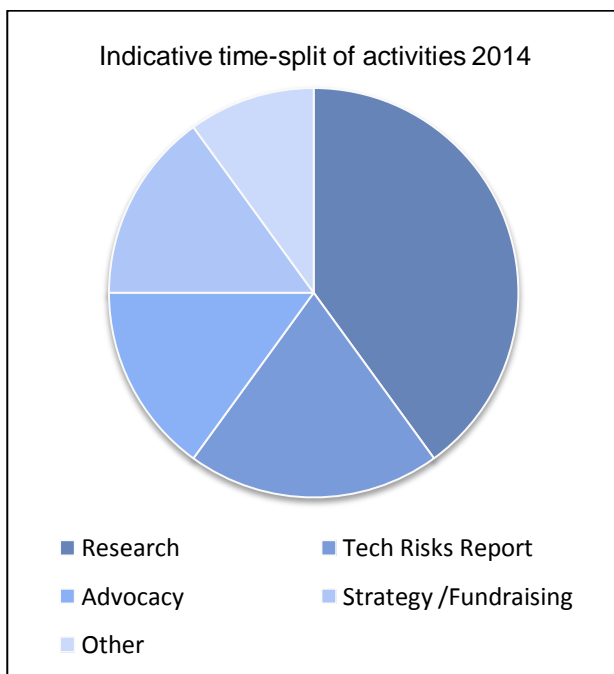
What are potential future avenues of research?

Although we have identified a prioritised list of activities for us to work on, there are many projects we would like to engage in but currently lack capacity to pursue. We are likely to address some of these in the future. In some cases we have a clear 'attack' on the question,

while in others we simply regard it as an important question of the sort we are interested in working on as an organisation.

These do not represent a current research plan for the Global Priorities Project. However, if a donor were to approach us with funds to support a particular topic, we would be open to prioritising it, although practical concerns could make it impossible at the time.

These include, but are not limited to, the following. We will curate and publish a longer list of questions (we have considered hundreds), and welcome discussion.



Policy evaluation in practice

- How well does the allocation of spending to prevention of topics on the UK Risk Register match their stated likelihood and impact?
- How could the methodology used by the Copenhagen Consensus, or similar organisations, be improved?
- What are the best ways to advocate for guidance on sensible scaling of numbers in politics (e.g., costs presented per user or per head of population)?

Policy evaluation in theory

- What are the best ways to compare outcomes which are presented in 'natural units' like years of schooling or DALYs averted? What is the best framework with a shared unit?
- Is there a natural way to extend QALY frameworks to include economic as well as health effects? Externalities of conditions for

families, carers, and communities? Broader measures of well-being? Population effects?

- What heuristics should we use to process new evidence to update our priors? How can we apply 'regression to the mean' heuristics to get more accurate policy predictions?

Historical lessons for long-term strategy

- How good have organisations been, historically, at identifying the great challenges of the coming century? The solutions? What made them more likely to be good?
- How good have organisations been at averting crises for humanity? What were the success factors?
- How good are organisations at affecting the far future deliberately? What are the success factors?
- What are the best ways to encourage multi-polar cooperation to address crises for humanity?
- Do crises for humanity tend to come in 'chunks' (decisive moments) or emerge from processes? What are the implications for risk mitigation strategy?

'Soon vs. Later' modelling

- What are the implications of the Soon vs. Later risk [strategy model](#) on topics like climate change? What are the best ways to estimate key parameters?
- What are the implications of the Soon vs. Later [funding model](#) for foundations allocating budgets? What are the best ways to estimate key parameters?
- Are opportunities for funding improving or getting worse in crucial fields? How much should we expect this to continue?

Effective altruism strategy

- Is effective altruism movement building the best near-term activity? On average vs on the margin? Under what conditions? How does it compare with other options?
- What discount rates should which sorts of organisations and effective altruist groups use in which situations?
- What are the best meta-charity activities?
- What is the relationship between global catastrophic risks and existential risk?
- What is the likely impact of a range of traditional altruistic activities on the far future?

What are our assets?

The Global Priorities Project is in an excellent position to improve prioritisation:

- We have access to a strong network of senior policy-makers, foundations, and researchers giving us remarkable opportunities for our research to impact decisions.
- We have produced novel world-class research on prioritisation accounting for uncertainty and time preference.
- We have exceptional staff with outstanding and uncommon skills.
- We have existing ties to effective altruist organisations and have access to their research.
- We have an extremely low cost structure. Our researchers are altruistically motivated and our overhead costs are shared with the rest of CEA, allowing us a current fully-loaded cost of ~£36k per researcher (compared with ~£100k through the university).

The team

Sebastian Farquhar - Director

Before joining the Global Priorities Project Sebastian worked at management strategy consultancy McKinsey & Co. with experience advising organisations on healthcare, public policy and innovation strategy. He was on the core team launching 80,000 Hours - a career impact consultancy for effective altruists. He has a Masters in Physics and Philosophy from the University of Oxford.

Dr. Owen Cotton-Barratt - Director of Research

Owen researches for the Global Priorities Project and is a research associate at the Future of Humanity Institute. He teaches at the University of Oxford and has research experience in pure mathematics, voting theory, artificial intelligence, and moral uncertainty.

In addition to the core team, a number of associated researchers and staff at the Centre for Effective Altruism and Future of Humanity Institute are collaborating closely with the Global Priorities Project, in particular.

Dr. Toby Ord – Research Associate

Toby is a Research Fellow at the Future of Humanity Institute working on population ethics. He is a leading expert in health economics and practical ethics. Toby founded Giving What We Can, which has raised pledges of more than \$350 million to the most effective charities and has received extensive coverage in the international media.

Dr. Niel Bowerman – Policy Associate

Niel has a PhD in physics from Oxford University. He was Climate Science Advisor to the President of the Maldives, a member of President Obama's Energy and Environment Policy Team, and the Executive Director of Climatico. Niel was awarded Oxford University's inaugural Vice-Chancellor's Civic Award.

Rob Wiblin - Research Associate

Rob has previously conducted charity evaluations for Giving What We Can, and policy analysis for the Australian Department of Innovation, Treasury and Productivity Commission. He is currently Executive Director of the Centre for Effective Altruism.

¹ E.g., the [Poverty Action Lab](#), [What Works Centres](#), the [Education Endowment Fund](#)

² As argued in the 2014 report "[Wellbeing and Policy](#)" chaired by Lord Gus O'Donnell, former head of the UK's civil service

³ For example, difference in discount rates yield very different policy suggestions after one or two generations. Alternatively, differences between median and mean absolute risk in catastrophe planning can span orders of magnitude

⁴ "[Innovation: managing risk, not avoiding it](#)"

⁵ For example the [detailed tool produced](#) by councils in Manchester or the consulting services offered by the [New Economics Foundation](#)