

Cause prioritization research

Katja Grace

Table of Contents

Our process.....	1
What is the problem?.....	2
What is being done?.....	3
What could be done?.....	7
Questions for further investigation	12
Accompanying documents	13
Notes.....	13

Our process

We investigated this topic because we believe that research directed at prioritizing philanthropic spending at the level of the cause may be highly valuable. We have the impression that very little of it is done, and that it may improve the value of philanthropic spending substantially.

We talked to eight people involved with cause prioritization research about their own work and views on the area:

- Alexander Berger of GiveWell Labs
- Robert Wiblin on the Copenhagen Consensus Center (he works for the Center for Effective Altruism but recently discussed their work with them at length)
- Seán Ó hÉigearthaigh on the Future of Humanity Institute and Center for the Study of Existential Risk
- Owen Cotton-Barratt on the Global Priorities Project
- Paul Christiano, an independent researcher who has advocated for more cause prioritization research
- Gordon Irlam of Back of the Envelope Guide to Philanthropy, and GRI Charitable Foundation

- Kris Putnam-Walkerly on Putnam Consulting Group
- Paul Penley on Excellence in Giving

These discussions are the basis for many of the more qualitative claims made in this report. We supplemented these interviews with desk research.

Note that many of the interviewees and the author are associated with the 'Effective Altruism movement', a recent social movement advocating the support of cost-effective charities and other altruistic activities, based on explicit reasoning. This movement consists mostly of young people, and is centered in Oxford, UK and the San Francisco Bay Area, US. These people's views may thus be more than usually correlated for social reasons, and so agreement between them should be interpreted accordingly.

This research included a search for further relevant groups, starting from a list of possibilities kindly provided by the Center for Effective Altruism. Our tentative list of related groups is available [here](#).

What is the problem?

There are many available avenues for improving the world, in terms of social value. There is also much uncertainty around the value of specific options. It seems plausible that some avenues much more effective than others, even among avenues that receive substantial funding¹. Research which successfully distinguishes the more promising of these avenues and redirects a small part of over \$300 billion² spent on philanthropy annually them could thus plausibly multiply the value of that funding by a large factor, and so produce a lot of value. For instance, if \$5M could produce recommendations that improved the impact of 0.1% of philanthropic funding by 5% for 10 years, this would be equivalent to \$150M of extra philanthropic spending, and so well worth the \$5M investment. These do not seem like implausible figures³.

Prioritization efforts to date have focused on prioritizing interventions or organizations with similar immediate goals to one another, for instance comparing the effectiveness of treating two different diseases in the developing world⁴. This narrow prioritization appears to often be accepted as best practice among philanthropists, and can be successful in redirecting funding⁵. Priorities among larger causes - such as developing world health versus biological research - has seemingly received less attention, and is much more uncertain. Yet similar arguments suggest that prioritizing causes at this level may also be important: there are likely large differences in value between causes, and we are fairly ignorant, so marginal understanding could multiply the value of sympathetic altruistic spending by a substantial factor.

Reasons to focus on causes over interventions or charities are that they are relatively ignored, there are plausibly larger differences in value between popular ones, information about causes is more useful than information about charities for larger donors⁶, and some decision about causes has to be made before developing the deep understanding of an area required to compare interventions or charities there well, as we see in the standard practice of foundations⁷.

Current methodologies for cause-level prioritization are often considered rudimentary. Among prioritization researchers from the Effective Altruism movement interviewed for this project, there was broad agreement that a large part of the value of prioritization research in the short to medium term will come from learning how to do such research, and building long run capacity for prioritization in other ways⁸. This suggests another reason to focus on the level of causes: knowledge gained about causes is likely to have longer term and more robust value than knowledge about particular interventions or charities⁹.

There are some reasons to think that not enough cause prioritization is done, from the perspective of total social welfare. Arguably very little is done, compared to efficiency-sensitive philanthropic spending¹⁰, given apparent difficulty. Many reasons have been given to expect this (see 'Further considerations for funding research'). On the other hand, the fact that so little cause prioritization appears to be done is also some negative signal about its worth, given that many people might be motivated to fund it.

As mentioned above, prioritization research appears to be in its infancy, and a lot of the value is in learning and developing the infrastructure for much more sophisticated prioritization in the future. The long term nature of this kind of output makes it likely to be underfunded from a social standpoint, as does the fact that it will be mostly useful for future people. That it involves diffuse benefits and concentrated costs is further reason to be unsurprised if prioritization is underfunded, though these facts do not set it apart from many other philanthropic causes¹¹.

For a more in-depth look at the promise of cause prioritization research, see Paul Christiano's draft, [The Case for Cause Prioritization as the Best Cause](#).

What is being done?

A small number of organizations work directly on prioritizing causes. Others sometimes produce prioritization between causes, though this is not their main project. Many more contribute to cause prioritization indirectly, for instance by producing relevant primary research, integrating primary research (e.g. in reviews), promoting cause-neutral giving, or funding according to cost-effectiveness advice.

Groups and people (with estimated budgets or sizes where available) doing substantial systematic cause prioritization¹² include the following:

Think tanks and other nonprofits searching for the highest priority causes (~\$3M)

The three listed here are probably most of this class.

- The Copenhagen Consensus Center (\$1-2M¹³) commissions top economists to prioritize interventions across a range of causes, and conduct outreach to promote their work.
- Givewell Labs (~\$1.2M¹⁴) is an early stage project seeking promising causes for foundation programs. It is highly involved with the foundation Good Ventures, so might also be well placed in the foundation research category.

- The Global Priorities Project (one researcher¹⁵) is a very new project of the Future of Humanity Institute and the Center for Effective Altruism. It has not produced substantial cause prioritization yet, but is intended to. It is experimenting with a variety of projects in the short term.

Policy researchers with other goals, using a common metric

Organizations trying to improve government spending within narrow areas sometimes use measures of cost-effectiveness which can be compared across different areas. Similarly, policy researchers in academic economics, think tanks and government sometimes produce research evaluating a specific policy, relative to some status quo. Both of these endeavors sometimes produce cost-benefit analyses of policies for instance, which estimate the total social value of the policy¹⁶. These can be compared between interventions from different areas if they use a relevant and comparable metric for value¹⁷, though if not intended to be used in that way, this may be misleading¹⁸. Below are two related efforts to influence government policies. They are part of the movement for 'Evidence Based Policy' which is probably more broadly relevant¹⁹. There are also many 'best practices' sites²⁰.

- Washington State Institute for Public Policy (15 people²¹) calculates benefits and costs for a range of Washington State public policy options²².
- The Pew-MacArthur Results First Initiative helps state governments to implement a cost-benefit approach to assessing the value of public programs²³. They provide a national database on program effectiveness, which allows calculation of costs and benefits to programs in combination with state-specific data²⁴.

Research by philanthropists generally doesn't include systematic evaluation of causes²⁵, however a small number of foundations have a strong stated interest in cost-effectiveness, and support projects arguably in multiple areas, so it is plausible that they do some systematic cause prioritization²⁶. These include for instance the Gates Foundation, the Hewlett Foundation²⁷, Good ventures²⁸, The Mulago Foundation²⁹, The Peery foundation³⁰, and Jasmine Social investments³¹. This research tends not to be public however, which limits influence and precludes other research building upon it. There are some efforts to increase transparency and improve these issues (e.g. GlassPockets, and The Center for Effective Philanthropy), and at least one example of a high degree of transparency:

- The Back of the Envelope Guide to Philanthropy (20-60h per year³²) publishes rough calculations evaluating cost-effectiveness of interventions across a wide range of causes. It guides the giving of GRI Charitable Foundation (to be \$200K/y in coming years³³), but is intended for outside use. It is conducted by GRI Charitable Foundation director Gordon Irlam in his free time.

Philanthropic consultants help nonprofits and philanthropists with many parts of their operations, including selecting focus areas³⁴. This can include systematic evaluation of opportunities across different causes³⁵. The size of this area is unclear. The following organizations have been interviewed in this project or by GiveWell; there are many others.

- The Bridgespan Group (\$33M³⁶)³⁷.
- Excellence in Giving (>6 people³⁸)
- Putnam Consulting Group (~1 person³⁹)

Many groups are relevant to cause prioritization because they produce important inputs to it, though they do not set out to prioritize causes themselves. The following are examples, emphasizing those that interviewees find important, and not excluding any known examples judged to be unambiguously important.

- Academic economists (17,230 people⁴⁰) spend some small fraction of their efforts researching questions that are especially relevant to cause prioritization, such as ‘what are the long term effects of education?’⁴¹ They also produce research relevant to evaluating many specific policies.⁴²
- The Center for High Impact Philanthropy (~15 people⁴³) publishes recommendations on interventions within causes, and other summary information intended for high net worth donors. They sometimes make recommendations across a number of causes at once⁴⁴, but there is no sign that these weight different causes systematically, rather than just listing good opportunities from different areas side by side.
- The Future of Humanity Institute (£700,000⁴⁵) researches questions relevant to ensuring humanity’s long term wellbeing, including for instance the importance of existential risk reduction.
- Producers and aggregators of primary research about developing world interventions, such as the Disease Control Priorities Project, Innovations for Poverty Action, and J-PAL
- Cause advocates and others using common metrics to promote single causes could in principle produce prioritization, if their procedure for producing an estimate is sufficiently comparable. QALYs and economic value seem to be used relatively often by different parties, sometimes for interventions in different causes.
- Independent researchers, for instance in the Effective Altruism community sometimes do small research projects answering questions intended to bear on cause prioritization⁴⁶.

Cause prioritization is part of a larger sector of efforts to direct philanthropic spending. This includes charity evaluators, philanthropic advisors, charity advertising, foundation research, think tanks, and cause advocates. This larger area may be crowded, even if cause prioritization is an uncrowded approach. This is relevant to the extent that different types of pressure on funding are substitutes.

Efficiency sensitive funding

Total giving to charitable organizations in the US was \$316.23 billion in 2012⁴⁷. Of that, \$47.44 billion was given by foundations, and \$18.97 billion by corporations. Global giving is probably

less than twice as large as US giving⁴⁸. Globally, official development assistance (ODA) - or aid originating from a set of relatively developed countries⁴⁹ - was around \$125 billion in 2012⁵⁰. Government spending is also often intended to create social value in broadly similar ways to philanthropy (e.g. through education, health, environmental wellbeing), and to some extent is driven by public perceptions about what will improve society, or more explicit evaluations of outcomes⁵¹. In the long run it is plausible that any of these would be influenced by research which changed public views on the value of an intervention. In the shorter term, there are a number of funders who are interested in cost-effectiveness.

Some funders who appear to be interested in cost-effectiveness and more than one cause are listed below. These foundations are commonly believed to be interested, and discuss their interest on their website. This is unlikely to be a reliable criterion, especially given that cost-effectiveness and similar ideals are becoming fashionable things to mention on websites, and common beliefs could easily originate from websites. See 'Questions for further exploration' for details on how one might better evaluate relevant funding. Also note that cause prioritization is likely to be more relevant to new foundations than existing ones, so the inference from large existing foundations who care about cost effectiveness is not as direct as one might imagine.

- Bill and Melinda Gates Foundation (\$40.2 billion endowment, \$3.4 billion granted in 2012)⁵²
- Good Ventures⁵³ (Granted \$3.9M in 2013, \$9.5M so far in 2014)⁵⁴
- The William and Flora Hewlett Foundation⁵⁵ (\$7.74 billion in assets, \$304 million in grants 2012)⁵⁶
- Mulago Foundation⁵⁷
- Peery Foundation⁵⁸
- Jasmine social investments⁵⁹
- Acumen Fund⁶⁰
- Giving What We Can members (\$250M pledged⁶¹)
- GRI Charitable Foundation (\$200K/y intended in coming years⁶²)

Some example outputs

- CCC research has been cited as a source of evidence in at least \$5 billion worth of aid spending decisions⁶³. For instance the UK Government invested \$4 billion in nutrition, citing CCC as their source for the high cost-effectiveness of direct nutrition interventions⁶⁴. It is presumably hard to get good evidence about the extent to which they changed those decisions.

- In 2013 GiveWell Labs claims to have moved \$965,000⁶⁵.
- Economists and other policy thinkers have produced analyses that appear to have improved educated discussion of the importance of different causes⁶⁶.
- The Pew-MacArthur Results First Initiative has produced a US national database of evidence on program effectiveness⁶⁷ and claim to have caused six US states so far to redirect \$38 million to more effective interventions⁶⁸.
- The Future of Humanity Institute's arguments regarding the overwhelming importance of the future and existential risk appear to have influenced many people's decisions, especially in the Effective Altruism movement⁶⁹.

Intangible outputs: learning and public discourse

Three interviewees said that a large fraction of the value of prioritization research in the short to medium term will come from learning and improving methodologies⁷⁰. If cause prioritization will ultimately be valuable, then early research may have substantial indirect value by determining which approaches are effective and providing a model for future research, even if early research itself has less immediate impact.

There are also likely benefits via improving discourse on these questions. GiveWell intends to do this, and appear to have some success⁷¹.

What could be done?

Improving the prioritization of philanthropic spending involves several activities other than prioritization research itself. This section will list plausible ways to contribute in this broader range of ways. Prioritization is informed by basic research on many topics, and aggregation and interpretation of that research. The success of prioritization depends on funders using its recommendations, so outreach or funding based on recommendations are other ways to contribute.

Relevant primary research

There appears to be a 'backlog' of existing academic work that could help to evaluate many philanthropic causes⁷². Sources such as academic economics research, DCP2, and field experts seem to be most useful. However this primary research is generally not created with the intention of informing philanthropic cause prioritization, so its overlap with what is needed for prioritization is limited. Gaps in research that are especially salient to interviewees include:

- Research relating to academically obscure causes, e.g. risks from artificial intelligence. Primary research tends to focus on politically salient causes, such as minimum wage and immigration laws. To forward this, one might support The Future of Humanity Institute⁷³.
- Research on long term and indirect effects: most empirical work naturally focuses on short term relatively direct benefits of interventions, such as lives saved. These are

substantially easier to measure than longer term effects. There is debate over both whether long run effects are highly relevant to cause prioritization or not, and whether they are prohibitively difficult to learn about⁷⁴.

- Modular, robust, well researched answers to small questions that are relevant to particular causes, based on aggregating empirical evidence.
- Organization of existing research: the relevant academic research that exists appears to be harder than necessary to find.

Prioritization research

Existing projects

Some ways to contribute through existing projects:

- The Copenhagen Consensus Center seeks funding, most recently for their post-MDG project. For that project, funding would probably do something like increase outreach for the research they have done, allowing them to avoid scaling back engagement with the media and officials⁷⁵.
- The Global Priorities Project welcomes funding, which it would spend on hiring more researchers.
- GiveWell Labs is very difficult to contribute to through funding; donations are unlikely to increase effort invested⁷⁶.

New projects

Some projects which could be run:

- Commission professional philanthropic advisors to produce overviews of causes. They appear to have relevant expertise and be willing to do such research, but it is rarely requested of them. This could be a quick way to try variants on the main products of the main prioritizing organizations, for the purpose of learning. It could also produce useful inputs to prioritization.
- Commission independent contractors to do any of a large number of small, modular projects. The EA community contains many capable and interested people, some of whom already do such research, on a paid or unpaid basis⁷⁷. Prioritization research seems relatively amenable to modular independent contributions, and relevant research inputs seem even more so.

Relevant products a project might have:

- Recommended causes (GiveWell Labs intends to produce this⁷⁸)
- Recommended interventions (the Copenhagen Consensus Center does this, as does WSIPP)⁷⁹.
- Cases for causes or interventions that can be compared

- Good frameworks for making comparisons: this is likely to be combined with making comparisons, but one may intentionally focus on producing a framework that others will want to use when evaluating opportunities. GPP has such a focus⁸⁰. This may produce value by causing others' promotion of their own causes to contribute to prioritization.
- Widely applicable higher level research e.g. how economic development in the developing world affects global economic trajectories, answers to methodological questions (e.g. GPP is looking into strategies for evaluating opportunities given a large degree of uncertainty about costs⁸¹).
- Information about causes that might help viewers to prioritize them, perhaps by a donor using their own values and judgement calls. For instance, the landscapes made by philanthropic advisors⁸² and shallow investigations made by GiveWell Labs⁸³.

Some salient project parameters:

- Whether to try to account for indirect effects: Everyone seems to agree that this is hard, but opinions vary on whether it is prohibitively so⁸⁴. Opinions also vary on whether such effects are an overwhelming part of the value of interventions, or a small one⁸⁵.
- Level of quantification: quantification brings rigor, transparency, and the ability to talk about judgements. On the other hand, it often results in leaving out vague considerations, and considerations where one does not have a good conscious model, and can give a false appearance of accuracy⁸⁶. Systematic cause prioritization sets itself apart from more traditional methods of selecting causes in part through a certain level of quantification. However among those practicing prioritization, there is some disagreement about how far one should go⁸⁷.
- Level of abstraction: GWL makes an effort to be grounded in concrete funding opportunities, while for instance GPP is happy to pursue abstract lines of research⁸⁸.
- Discovery vs. sorting good causes: one could create value by looking for good new causes, or by evaluating causes that appear to be reasonably good.
- Relationship with existing projects: overlap vs broader coverage It is unclear whether GiveWell prefers other researchers doing similar projects were to evaluate the same causes they have evaluated, providing a check, or to evaluate causes GWL hasn't looked at⁸⁹. Gordon Irlam thinks newcomers would do best to research its current top recommendations more thoroughly, to produce a more convincing case for them⁹⁰.
- Intended audiences: current projects are generally directed at philanthropists and government policymakers. CCC's current project is directed at people participating in the post-2015 Millennium Development Goals process of the United Nations⁹¹. Other possible audiences include charities, people choosing careers, and academics.

- Object level arguments vs. other indicators of quality – one could evaluate causes based on evaluation of the evidence. Alternatively, one can collate indicators such as expert opinion, and room for more funding.
- Values: to prioritize interventions, it is necessary to assume some values. There is a question then of which ones to use, which is made more difficult if you are aiming to influence others' spending. Plausible possibilities include a metric of long term aggregate social welfare, your own values, the likely values of the target audience, or to try to provide recommendations which are robust to variations among popular choices of values.

Further considerations for funding research

There appears to be a risk in this kind of research of outputs not being used, especially if the research is done outside or not in connection with the parties who might use it. It seems valuable to have a plan for how the work will affect specific party's decisions⁹². It may be even be more useful to encourage philanthropists to either use existing resources or to do such research themselves.

Reliable comparison research is often considered difficult, and prioritizing causes appears to be harder than prioritizing narrow interventions. It is harder to find common metrics for more different activities, and the search for common metrics leads higher into abstractions and deeper toward fundamental values⁹³. Relatedly, prioritizing causes involves either knowing more or making more assumptions about long run effects than does narrower prioritization, because with similar interventions many effects can be factored out.

Given that a large amount of value seems to come from learning at the moment, it may be good to design projects to capture and use this value effectively. For instance, projects that are flexible, have fast feedback loops, and produce some output from the learning such that others may benefit from it⁹⁴.

One should also ideally have some idea of why there appears to be little work on this issue, and how one's effort makes sense in light of this explanation. For instance, if your explanation is that philanthropists will not use such advice, then your project might do better to address that disinterest (or do something else entirely), than to produce more advice that will not be used.

Plausible partial explanations for low provision of such research include:

- Philanthropists generally do not care enough about social value per se to do costly research, though they will respond to information if it is readily available.
- Philanthropists have specific other interests or pressures upon their spending, and so information about social value is not very relevant to them

- People perceive (correctly or otherwise) that there are more efficient ways to improve value from philanthropy than by evaluating causes
- Philanthropists generally choose issues of interest long before they have the resources that would warrant a huge investment in evaluating causes
- People care about the good they do personally, not the total good that is done, so they are unwilling to sacrifice their own funds to produce information that will allow others to give better.
- It is hard for philanthropists to trust outsiders, because many groups incessantly vie for their money, so it is very hard to advise them⁹⁵
- Philanthropists do such research at a locally optimal level, but do not share the research, so the optimal level is below what it would be if such research were reused.
- Such research would have distributed gains and concentrated costs, so provision is classically difficult to coordinate.
- It is hard to reuse such research, because different funders have different values, and the objective situation changes rapidly.
- Such research is very hard

Other research-related initiatives

- Encourage proponents of interventions or organizations to use a common, meaningful and comparable format in their requests for support. e.g. collaborate with a funder to require that requests include an estimate of QALY/\$.
- Encourage cause prioritization research to taken up more in academia, and seen as more academically respectable. e.g. fund post-docs.
- Coordinate interested amateurs to do in-depth work together, or try out this kind of work on a small scale. e.g. choose a larger project together, such as evaluating education improvement in the developing world, and break it up into parts, so that for instance one person investigates the effect of education on fertility, and another the effect on income, etc.
- Improve coordination and sharing among foundation researchers. There casually appears to be a lack of this, perhaps making foundation research more difficult and less useful than it could be. This could involve collaborating with projects like [GlassPockets](#), [IssueLab](#), and [The Center for Effective Philanthropy](#).

Outreach

Value from prioritization research is highly contingent on it being used by funders. Some feel that outreach is of primary importance at the moment (e.g. Gordon Irlam). Others argue that most projects are too much in their infancy, and prioritization research will be much easier to sell when it has better products for funders to use (e.g. Paul Christiano and Owen Cotton-Barratt)⁹⁶.

Most funders do not appear to be interested in systematic evaluation of opportunities from an aggregative social welfare perspective. Yet the idea of cost-effective philanthropy is gaining traction in the public sphere, so increasing interest among funders may be relatively straightforward (though perhaps also inevitable).

- Pragmatic outreach: outreach has focused on the merits of philosophical ideas around explicit cost-effectiveness optimization, but there may be value in addressing more concrete issues with funders, such as how to go about prioritizing better, and how they can profitably share research. Collaboration with organizations trying to improve philanthropy in similar ways might make sense here, e.g. The Foundation Center, which runs glasspockets.org, and issuelab.org, aimed at increasing transparency and information sharing in the philanthropic sector respectively.
- Deliver workshops at philanthropy conferences, e.g. Gordon Irlam suggests the [EDGE](#) conference in Berkeley or those listed by the [Chronicle of Philanthropy](#)⁹⁷.
- Encourage people who care about cause prioritization into foundation research work⁹⁸.
- Coordinate with similar movements, such as the evidence based policy one.

Questions for further investigation

- What is the landscape like for smaller areas of cause prioritization? This has been a very broad investigation. A good next step might be to do a similar investigation at a smaller scale, in just one area of this larger cause. For instance, in promoting information sharing and transparency in foundations, economic research on long-run consequences, or comparison methodology work.
- How much sensitive funding is there really? An important step in looking to fund such research would be accurately assessing what kind of research would move how much money. Such research might address questions such as: Has the group chosen a cause? If so, is there evidence they did this based on evidence? Do their high level funding decisions in general depend on research? Do their decisions depend on outside research? Are they transparent enough that we can be confident that their self-description matches their decision-making?
- How much is spent on trying to shift philanthropic spending in total? (including in many other directions than toward cost-effectiveness) This is relevant because cause prioritization efforts compete with all such work.

- How relevant is work by related groups? For instance, philanthropic consultants, organizations for philanthropic transparency, organizations for evidence based policy and for-profit efforts to find valuable ventures. Should people interested in cause prioritization be interacting with them?
- How much does it cost to make a policy-relevant recommendation, and how valuable is such a recommendation? This work has included some suggestive figures, but a more detailed assessment would be useful.
- What are pressing questions in cause prioritization? This would include consideration of who might be moved by better information, and what their uncertainties are.

Accompanying documents

Conversation with Rob Wiblin about the Copenhagen Consensus Center

Conversation with Alexander Berger about GiveWell Labs

Conversation with Gordon Irlam about Back of the Envelope Guide to Philanthropy

Conversation with Paul Christiano about Cause Prioritization

Conversation with Owen Cotton-Barratt about the Global Priorities Project

Conversation with Seán Ó hÉigeartaigh about the Future of Humanity Institute and the Center for the Study of Existential Risks

Conversation with Kris Putnam-Walkerly about Putnam Consulting

Conversation with Paul Penley about Excellence in Giving

I'm grateful to Nick Beckstead, Paul Christiano, Ben Todd, Matt Wage, Anna Altea, the Center for Effective Altruism, and all of the interviewees for their help with this project.

Notes

¹ This is commonly believed among people in the Effective Altruism movement. For instance, GiveWell write "Our default assumption, or prior, is that a charity – at least in one of the areas we've studied most, U.S. equality of opportunity or international aid – is falling far short of what it promises donors, and very likely failing to accomplish much of anything (or even doing harm)." (see <http://blog.givewell.org/2009/12/05/a-conflict-of-bayesian-priors/>). Paul Christiano would not be surprised

to discover that the most effective philanthropy in the future was ten or a hundred times more effective than contemporary philanthropy (see conversation). These intuitions are partly driven by large discrepancies in the direct health effects per dollar of interventions that have been studied (see http://www.givingwhatwecan.org/sites/givingwhatwecan.org/files/attachments/moral_imperative.pdf, and note that the cost-effectiveness of interventions overall are also influenced by the size of indirect effects and differences in execution among different projects), along with pessimism about the pattern of funding being better among more difficult to study interventions.

² The US gave \$316.23 billion to charitable organizations in 2012.

(<http://www.charitynavigator.org/index.cfm?bay=content.view&cpid=42%20-%20UoJFr74o59A#.U1FdleZdVnE>). It is unclear whether this includes activities by foundations who operate their own programs ('private operating foundations'), or how much is spent by such foundations. Total spending worldwide seems likely to be less than twice this amount. The US represents around 22% of global GDP, however appears to have substantially higher rates of philanthropy per dollar of GDP (2%) than average. For instance, China represents 12% of global GDP, and apparently gave around \$5 billion in a recent year (or 0.06%)

(<http://www.knowledgeatwharton.com.cn/index.cfm?fa=printArticle&articleID=2337&languageid=1>). The UK spent \$15.6 billion of its \$2.4T GDP, for 0.6%

(<https://www.cafonline.org/PDF/UKGiving2012Summary.pdf>).

³ This would be roughly three times what is spent per year on the Copenhagen Consensus Center, for influence over funding worth one tenth of the Gates Foundation.

⁴ For instance, WHO-CHOICE, Giving What We Can, and until recently GiveWell, have focused strongly on finding good charities or interventions in developing world health. Philanthropedia recommends organizations across a range of causes, but only in comparison with other organizations in the same cause. Philanthropic advisory firms generally search for good interventions within a narrow range of causes (see footnote 25), and it appears on casual review of websites that cost-effectiveness minded foundations generally choose a cause, and then search for cost-effective interventions within it. For the range of related organizations we are aware of, see [this list](#).

⁵ e.g. GiveWell has moved around \$30M to its top recommendations

(<http://www.givewell.org/about/impact>).

⁶ According to Alexander Berger (see conversation).

⁷ Suggested by conversation with Paul Christiano. See his answer to the question “Why is it better to evaluate causes than interventions or charities?”

⁸ From conversation with Alexander Berger: “GiveWell Labs is prioritizing learning value and diversification at the moment, and not aiming to make decisions about cause priorities once and for all. Alexander would guess that the impact of GiveWell Labs' current efforts is divided roughly equally between immediate useful research output and the value of trying this project and seeing how it goes.”.

From conversation with Owen Cotton-Barrett: “There are different routes to impact with cause prioritization. Owen thinks a major route to impact is through laying bricks of prioritisation methodology. This will help people in the future to do better prioritisation (and could be better than anything we manage in the near-term future)”.

Paul Christiano says, “This judgment is based on optimism about how much money could potentially be directed by the kind of case for impact which we could conceivably construct, but also the belief that there is a good chance that over the very long-term we can hope that the philanthropic community will be radically better-informed and more impactful (think many times) than it currently is. If that's the case, then it seems likely that a primary output of modern philanthropy is moving towards that point. This is not so much a story about quickly finding insights that let you find a particular opportunity that is twice as effective, and more a story of accumulating a body of expertise and information that has a very large payoff over the longer-term.”

Paul also says, “...So there is a lot to figure out about how to go about the figuring out, and I would imagine that the primary impact of early efforts will be understanding what settings of those parameters are productive and accumulating expertise about how to attack the problem.”

See conversations with Paul Christiano, Alexander Berger and Owen Cotton-Barratt.

⁹ Paul Christiano says “... there are just too many charities to do this analysis for many of them, and the landscape is changing over time (this is also true for interventions though to a lesser extent). If you want to contribute to a useful collective understanding, information about these broader issues is just more broadly and robustly applicable.” See conversation with Paul Christiano.

¹⁰ Probably less than \$10M/year is spent on intentional cause prioritization (see ‘What is being done’ below), though more is spent on other relevant activities. The scale of funding which might be sensitive to cause prioritization research is even less clear, but probably between hundreds of millions and ten billion

philanthropic dollars per year. See 'Efficiency sensitive funding' below for an idea of how much money may be moved by good research.

¹¹ See conversation with Owen Cotton-Barratt.

¹² Cause prioritization here refers to comparison of activities across different 'causes'. 'Causes' can be defined as GiveWell does, as 'a particular set of problems, or opportunities, such that the people and organizations working on them are likely to interact with each other, and such that evaluating many of these people and organizations requires knowledge of overlapping subjects' (<http://blog.givewell.org/2013/05/30/refining-the-goals-of-givewell-labs/>). Here 'cause prioritization' also includes evaluation of narrow interventions rather than whole causes, where the interventions span multiple causes (e.g. comparing reforestation for climate change prevention to investing in advocacy around a specific immigration law). This list is restricted to 'systematic' cause prioritization because everyone who spends money altruistically implicitly prioritizes causes. We are interested in those using more sophisticated methods. This list is also restricted to efforts to evaluate causes on their social benefits, rather than on other quality.

¹³ See notes on the Copenhagen Consensus Center

¹⁴ GiveWell Labs appears to be very roughly half of GiveWell, whose annual budget is \$2.5M (p4, <http://files.givewell.org/files/ClearFund/GiveWell%202012%20Audited%20Financial%20Statement.pdf>)

¹⁵ See conversation with Owen Cotton-Barratt.

¹⁶ For example:

Wang, Samuel J., Blackford Middleton, Lisa A. Prosser, Christiana G. Bardon, Cynthia D. Spurr, Patricia J. Carchidi, Anne F. Kittler, et al. "A Cost-Benefit Analysis of Electronic Medical Records in Primary Care." *The American Journal of Medicine* 114, no. 5 (April 1, 2003): 397–403. doi:10.1016/S0002-9343(03)00057-3.

Clemens, Michael A. (2011) "Economics and Emigration: Trillion-Dollar Bills on the Sidewalk?" *Journal of Economic Perspectives* 25, no. 3 (August 2011): 83–106. Available at <http://pubs.aeaweb.org/doi/pdfplus/10.1257/jep.25.3.83>

Brown, P., and A. Daigneault. "Cost–benefit Analysis of Managing the Invasive African Tulip Tree (*Spathodea Campanulata*) in the Pacific." *Environmental Science & Policy* 39 (May 2014): 65–76. doi:10.1016/j.envsci.2014.02.004.

¹⁷ For instance, <http://www.wsipp.wa.gov/BenefitCost> is evidently designed to compare interventions within causes. However we can infer from it that legislators can get around \$63 in social value per dollar spent on the best program in childhood education, whereas less than \$15 per dollar spent on any program in child welfare.

¹⁸ If for instance cost-benefit analyses of poverty reduction strategies are only intended to be compared against one another, any effect that always occurs proportionally in such strategies can be factored out (e.g. if every life you save brings about some small environmental costs), but this will make comparisons with a different cause inaccurate.

¹⁹ See a summary of some activities at http://opinionator.blogs.nytimes.com/2012/05/30/worthy-of-government-funding-prove-it/?_php=true&_type=blogs&_r=0 and the Center for Evidence Based Policy at <http://www.ohsu.edu/xd/research/centers-institutes/evidence-based-policy-center/>

²⁰ Here is a list of best practices sites: <http://ctb.ku.edu/en/databases-best-practices>

²¹ <http://www.wsipp.wa.gov/About/Staff>

²² Their benefit-cost results to date are in the categories of in justice, child welfare, education, mental health, general prevention, substance abuse, public health and housing, and can be seen at <http://www.wsipp.wa.gov/BenefitCost>.

²³ <http://www.pewstates.org/projects/pew-macarthur-results-first-initiative-328069>

²⁴ <http://www.pewstates.org/research/data-visualizations/the-pew-macarthur-results-first-approach-85899527140>

²⁵ In *Strategic Cause Selection*, Holden Karnofsky of GiveWell says 'Our picture of how most major foundations work is as follows: First, broad program areas or "causes" – such as "U.S. education" and "environment" – are chosen. This step is almost entirely "from the heart" – no systematic review is conducted, but rather the philanthropist (or foundation President) chooses areas s/he is passionate about'. Interviews with the philanthropic advisory firm representatives gave a similar picture for most

cases, though sometimes philanthropists select a region rather than a cause to narrow their search, and then consultants look at interventions in several pressing issues in that region.

²⁶ It seems difficult to find public discussion of the reasoning which these foundations use to pick focus areas beyond a very cursory analysis, with the exception of Good Ventures which provides such discussion but has not yet chosen a focus area.

²⁷ http://www.hewlett.org/uploads/documents/Outcome_Focused_Grantmaking_March_2012.pdf

²⁸ <http://www.goodventures.org/about-us/vision-and-values>

²⁹ The Mulago Foundation restricts itself to developing world issues, but within that category it funds relatively different types of intervention. It determines funding on a number of criteria, including a 'cost per key impact' metric, which it calculates for all organizations it funds. It is unclear how comparisons are made for different types of key impact.

³⁰ <http://www.peeryfoundation.org/approach/criteria>

³¹ <http://www.jasmine.org.nz/our-approach/>

³² Using Gordon Irlam's estimate of 200-600 hours total over the past ten years (see conversation).

³³ See conversation with Gordon, from 'Gordon also runs the..'

³⁴ See an interview between GiveWell and Bridgespan representatives at [http://files.givewell.org/files/conversations/Bridgespan%207-22-13_%20\(public\).pdf](http://files.givewell.org/files/conversations/Bridgespan%207-22-13_%20(public).pdf), and conversations with Kris Putnam-Walkerly and Paul Penley. Sometimes philanthropists select a region rather than a cause to narrow their search, and then consultants look at interventions in several pressing issues in that region.

³⁵ See conversations with Kris Putnam-Walkerly and Paul Penley.

³⁶ The Bridgespan Group Annual Report 2012, <http://www.bridgespan.org/getmedia/5d7a066b-1411-44dd-86f4-9b1e71c8f608/The-Bridgespan-Group-Annual-Report-2012.pdf.aspx>, p8

³⁷ See [http://files.givewell.org/files/conversations/Bridgespan%207-22-13_%20\(public\).pdf](http://files.givewell.org/files/conversations/Bridgespan%207-22-13_%20(public).pdf) for more details on their approach.

³⁸ <http://excellenceingiving.com/leadership/>

³⁹ This is strongly suggested by <http://putnam-consulting.com/about/kriss-bio/>

⁴⁰ There were 17,230 economists in the US in May 2013, each paid on average \$101,450 per year according to the US Bureau of Labor Statistics (<http://www.bls.gov/oes/current/oes193011.htm>). This gives \$1.7 billion in economist wages.

⁴¹ e.g. Mincer, Jacob, and Linda S. Leighton. Effect of Minimum Wages on Human Capital Formation. Working Paper. National Bureau of Economic Research, February 1980. <http://www.nber.org/papers/w0441>.

⁴² For example:

Anderson, S. & Coate, S. (2005), Market Provision of Broadcasting: A Welfare Analysis, Review of Economic Studies, available at <http://restud.oxfordjournals.org/content/72/4/947.short>,

Levaggi, R. (2014) Pricing schemes for new drugs: A welfare analysis, Social Science & Medicine, Volume 102, February 2014, Pages 69-73, available at <http://www.sciencedirect.com/science/article/pii/S0277953613006606>

Madestam, Andreas, Daniel Shoag, Stan Veuger, and David Yanagizawa-Drott. "Do Political Protests Matter? Evidence from the Tea Party Movement". *The Quarterly Journal of Economics* 128, no. 4 (November 1, 2013): 1633–85. doi:10.1093/qje/qjt021.

⁴³ According to <http://www.impact.upenn.edu/about/whoweare>, the project involves 15 people, though their level of involvement is unclear.

⁴⁴ e.g. http://www.impact.upenn.edu/international-issues/reports/category/holiday_giving

⁴⁵ See conversation with Seán Ó hÉigartaigh.

⁴⁶ For instance, Carl Shulman (<http://reflectivedisequilibrium.blogspot.com/search/label/cause%20prioritization>), and Nick Beckstead (<https://docs.google.com/viewer?a=v&pid=sites&srcid=ZGVmYXVsdGRvbWFpbnxuYmVja3N0ZWFKfGd4OjExNDJZTcwNjMxMzRmZGE>). Leverage research contains people endeavoring to find top causes (private conversation with Geoff Anders). They are focusing on techniques for finding very high value interventions that are not yet underway, rather than evaluating existing causes.

⁴⁷ <http://www.charitynavigator.org/index.cfm?bay=content.view&cpid=42%20-%20.UoJFr74o59A#.U1LXueZdVnF>

⁴⁸ See footnote 2.

⁴⁹ <http://www.oecd.org/dac/stats/totaldacflowsataglace.htm>, though note that other sources disagreed by up to a factor of three upwards. Giving countries can be seen at <http://www.oecd.org/dac/dacmembers.htm>

⁵⁰ <http://www.oecd.org/dac/stats/aidtopoorcountriesslipsfurtherasgovernmentstightenbudgets.htm>

⁵¹ A study by the Pew-MacArthur Results First Initiative found that all US states did some kind of cost-benefit analysis between 2008 and 2011, however only 36 of them used information from the analyses in policy decisions. See <http://www.pewstates.org/research/reports/states-use-of-cost-benefit-analysis-85899490452>

⁵² <http://www.gatesfoundation.org/who-we-are/general-information/foundation-factsheet>

⁵³ <http://www.goodventures.org/our-portfolio/grantmaking-approach>

⁵⁴ <http://www.goodventures.org/our-portfolio/grants-database>

⁵⁵ <http://www.hewlett.org/programs/effective-philanthropy-group>

⁵⁶ <http://www.hewlett.org/about-us>

⁵⁷ <http://www.mulagofoundation.org/how-we-fund>. The Mulago Foundation restricts itself to developing world issues, but within that category it funds relatively different types of intervention. It determines funding on a number of criteria, including a 'cost per key impact' metric, which it calculates for all organizations it funds. It is unclear how comparisons are made for different types of key impact.

⁵⁸ <http://www.peeryfoundation.org/approach/criteria>

⁵⁹ <http://www.jasmine.org.nz/our-approach/>

⁶⁰ <http://acumen.org/investments/investment-model/> mentions that they filter projects for high impact. Also their overall project is to build socially useful for-profit ventures in the developing world, which naturally entails a higher than usual degree of attention to outcomes (see <http://acumen.org/ideas/patient-capital/>).

⁶¹ <http://www.givingwhatwecan.org/about-us>

⁶² See conversation with Gordon, from ‘Gordon also runs the..’

⁶³ See

<http://www.givingwhatwecan.org/sites/givingwhatwecan.org/files/reports/CCC%20MDG%20Report.pdf> p9 for a longer list of examples of CCC’s impact.

⁶⁴ See citation vii in

https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/248760/Endorserscompact_update7_10_2013.pdf.

⁶⁵ GiveWell Metrics Report – 2013 Annual Review (2014),

<http://files.givewell.org/files/images/GiveWell%20Metrics%20Report%20%E2%80%93%202013.pdf>, p2

⁶⁶ For instance Clemens <http://pubs.aeaweb.org/doi/pdfplus/10.1257/jep.25.3.83> lists 25 other papers which estimate gains from elimination of barriers to labor mobility, capital merchandise trade in terms of efficiency gains as a percentage of world GDP. His analysis appears to have been influential; see <http://openborders.info/double-world-gdp/> and <http://econfaculty.gmu.edu/bcapan/pdfs/whyimmigration.pdf>

⁶⁷ <http://www.pewstates.org/research/data-visualizations/the-pew-macarthur-results-first-approach-85899527140>

⁶⁸ <http://www.pewstates.org/research/analysis/achieving-success-with-the-pew-macarthur-results-first-initiative-85899538638>

⁶⁹ Paul Christiano, a member of the EA movement, says “Many EAs have been influenced by arguments regarding astronomical waste, existential risk, shaping the character of future generations, and impacts of AI in particular. To the extent that we actually have aggregative utilitarian values I think these are hugely important considerations and that calling them to attention and achieving increased clarity on them has had a positive impact on decisions (e.g. stuff has been funded that is good and would not have been otherwise) and is an important step towards understanding what is going on. I think most of the positive impact of these lines of argument will wait until they have been clarified further and worked out more robustly”. See conversation with Paul Christiano.

For more on FHI’s recent accomplishments, see conversation with Seán Ó hÉigartaigh, p2.

⁷⁰ Alexander Berger estimates around half of GWL benefits are in trying this project and seeing how it goes (see conversation with him). Paul Christiano and Owen Cotton-Barratt express similar sentiments (see discussions).

⁷¹ GiveWell, on GiveWell Labs: 'Our goal is not to apply the "evidence-backed/thoroughly-vetted/underfunded" framework to causes where it may be out of place, but rather to improve the level of transparency and public dialogue around how to make the most of one's giving(<http://www.givewell.org/givewell-labs>). For evidence of success, see http://en.wikipedia.org/wiki/GiveWell#Media_and_blog_coverage,

⁷² See conversation with Owen Cotton-Barratt.

⁷³ FHI welcomes donations: smaller amounts of funding would be spent on thesis prize competitions and other research field-encouraging activities, flying in journalists, and building reserves. Larger sums would be spent on building up core operations staff, funding current researchers, and hiring researchers in new areas. See conversation with Seán Ó hÉigearthaigh.

⁷⁴ e.g. From conversation with Owen Cotton-Barrett: "It has been suggested by others that research on long run consequences of policies is prohibitively difficult. Owen believes that improving our predictions in expectation about these long run consequences is hard but feasible. This is partly because our predictions are currently fairly bad."

Paul Christiano also finds long term considerations worth investigating (see <http://paulfchristiano.com/philanthropy/>).

From conversation with Alexander Berger: "Answering high level questions such as 'how good is economic growth?' doesn't seem very decision relevant in most cases. This is largely because these issues are hard to pin down, rather than because they are unlikely to make a large difference to evaluations if we could pin them down, though Alexander is also doubtful that they would make a large difference. For instance, Alexander doesn't expect indirect effects of interventions to be large relative to immediate effects, while Holden Karnofsky (co-executive director of GiveWell) does, but their views on this do not seem to play a big role in their disagreements over what GiveWell Labs should prioritize.

⁷⁵ This was true when Robert Wiblin spoke to them in 2013, and CCC is still advertising this project above others on their website and seeking funding for it (<http://www.copenhagenconsensus.com/make-donation>), however it is unclear whether the marginal activities on the project are still the same.

⁷⁶ See conversation with Alexander Berger.

⁷⁷ Leverage research is open to discussing earmarked donations for a cause prioritization project, for instance.

⁷⁸ <http://www.givewell.org/givewell-labs>

⁷⁹ See <http://www.copenhagenconsensus.com/copenhagen-consensus-2012/outcome> and <http://www.wsipp.wa.gov/BenefitCost>.

⁸⁰ See conversation with Owen Cotton-Barratt.

⁸¹ For more on this, see conversation with Owen Cotton-Barratt.

⁸² [http://files.givewell.org/files/conversations/Bridgespan%207-22-13_%20\(public\).pdf](http://files.givewell.org/files/conversations/Bridgespan%207-22-13_%20(public).pdf)

⁸³ <http://www.givewell.org/shallow>

⁸⁴ See footnote 74.

⁸⁵ Owen Cotton-Barratt suspects they are large (see conversation with him), Gordon finds them ‘because of his selection process, inherently less important’ (see conversation with him), Alexander and Holden of GiveWell Labs disagree (see footnote 74), and Paul thinks they are important (see <http://www.jefftk.com/p/flow-through-effects-conversation>).

⁸⁶ See conversation with Owen Cotton-Barratt.

⁸⁷ For instance, Gordon Irlam is in favor of ‘quantifying everything’ (see discussion), while GiveWell has argued against using quantification or taking quantified evaluations seriously (e.g. see <http://blog.givewell.org/2013/08/08/passive-vs-rational-vs-quantified/>, <http://www.givewell.org/international/technical/criteria/cost-effectiveness>, <http://blog.givewell.org/2011/08/18/why-we-cant-take-expected-value-estimates-literally-even-when-theyre-unbiased/>).

⁸⁸ Alexander Berger suggests cause prioritization work can become fairly abstract, which he suspects makes it less useful to GiveWell Labs (see conversation). GWL tries to keep grounded by looking for concrete funding opportunities. Owen prefers a back and forth between concrete questions and abstract issues they bring up (see conversation with him).

⁸⁹ See conversation with Alexander Berger.

⁹⁰ See conversation with Gordon Irlam

⁹¹ <http://www.copenhagenconsensus.com/post-2015-consensus>

⁹² Alexander recommends this (see conversation).

⁹³ Across a wide range of values, people can agree that saving two people is better than saving one. Whereas many differences in values might affect whether saving five polar bears is better than saving one person.

⁹⁴ GiveWell exemplifies this well. See <http://www.givewell.org/about/transparency>.

⁹⁵ See conversation with Gordon Irlam.

⁹⁶ See conversation with Paul Christiano from 'I think that the number of people who might care', and conversation with Owen Cotton-Barratt from 'in response to the suggestion'.

⁹⁷ See conversation with Gordon Irlam.

⁹⁸ Suggested by Owen Cotton-Barratt; see conversation notes.