EXPLORATORY ALTRUISM

PRIMARY AUTHORS: S. HILTON & V. AGARWALLA

Review: J. Savoie

JULY 2021

RECOMMENDED

Background to this research

Charity Entrepreneurship helps start high-impact nonprofits through researching promising interventions and running an annual Incubation Program to launch the best of them.

In 2021, we researched top interventions that could be implemented in the **effective altruism meta** space. Meta work helps individuals or other organizations accomplish their goals. For example, effective altruism (EA) meta might look at the best ways to help the EA movement accomplish its goal of maximizing good through reason and evidence. Charity Entrepreneurship is itself an EA meta charity, doing good through helping other charities get started instead of directly implementing interventions. This post discusses why we selected EA meta as a cause area.

Our 2021 recommendations in the EA meta space are training, exploratory altruism, and earning to give +. To arrive at our three recommendations we used a many weak arguments approach, including synthesizing views from across the EA community (with a survey of forty EA community leaders), and assessing our priors, cross-applicable data, crucial considerations and a need for flexibility. Our research on these three top ideas involved among other tools cost-effectiveness analysis, in-depth expert interviews, theory of change design, and informed consideration.

Please see the annex for further discussion of our EA meta research methodology, and refer to our detailed research process for a full discussion of the core ideas and methodology used for other causes.

This writeup assumes the reader has some background knowledge of the effective altruism (EA) movement. Readers not familiar with it can learn more about it here.

Acknowledgments

Thank you to Peter Wildeford for review and feedback, and to Karolina Sarek, Joey Savoie, Antonia Shann, and Urszula Zarosa for their contributions to this report. We are also grateful to all the anonymous experts who so kindly gave their time and thoughts in interviews and surveys.

Table of contents

1 What is Exploratory Altruism?

Introduction

Effective altruism (EA) is the question of how to do the most good. As the movement is relatively young, many areas and causes still have not been examined in depth. Given the amount of unexplored ground, it is likely that multiple highly promising causes are not yet on the EA radar.

A unique aspect of the EA movement is its openness to new ideas and paths to impact. There is huge potential for more exploratory cause prioritization research to provide value to the EA community. There may be a cause X that is highly impactful yet neglected by altruists.. There may be people who cannot achieve their full potential for impact due to a lack of research into which causes match their skills and goals.

A new organization could focus on exploring the case for new cause areas, and (to a lesser extent) on conducting cross-cutting research relevant to cause prioritization.

What do we mean by exploration?

To develop this idea, we have broken down the range of exploratory cause prioritization research into the following rough categories:

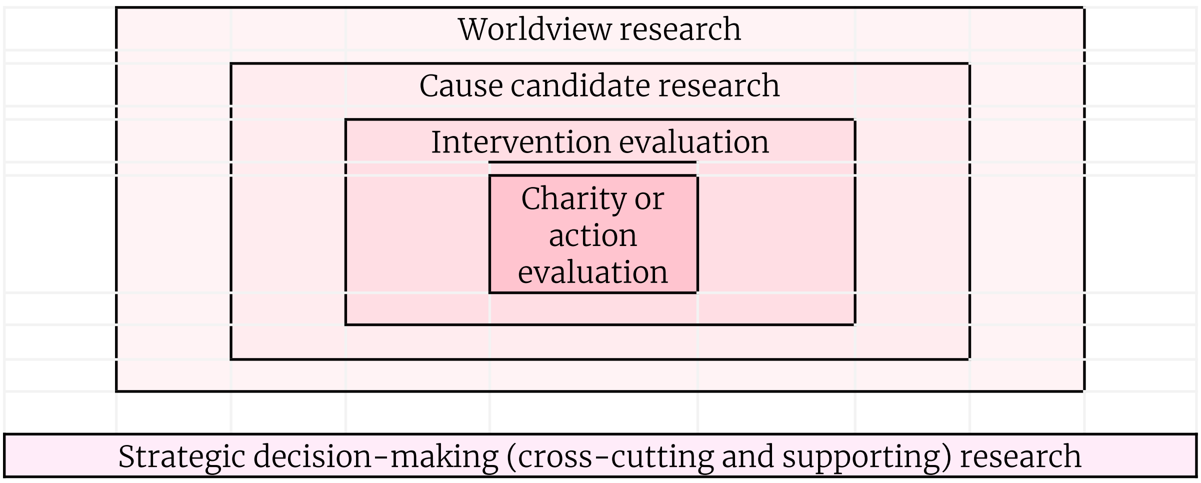

We would be most excited by an organization that focuses on **cause candidate research**, although we also think there are likely to be topics within **strategic decision-making** that could support cause research and might be particularly useful for this organisation to look at.

These categories can be defined as follows.

Worldview research:
Broad scope research into what matters in the world. This could involve researching ethics, worldview-relevant empirics (how to make worldview-level decisions under uncertainty, e.g. exploring concepts like cluelessness, cluster thinking, moral parliaments, etc.), and worldview-relevant crucial considerations (e.g. longtermism vs neartermism, animal sentience, the nature of human emotions).

Cause candidate research
Evaluating various specific causes. This could be topics such as great power conflict, voting reform, women's empowerment, economic growth and so on. At the simplest level, research could involve making the strongest possible case (i.e. a steelman) for each cause. A more complicated approach would be to conduct comparative research that examines different causes from different worldviews or ethical perspectives.

Intervention evaluation
Detailed investigations of specific ways of doing good. This could work on a few levels, such as:

- High-level (within a cause) – e.g. is reducing lead exposure a good way of improving global health?
- Granular – e.g. which lead exposure interventions are the most promising?

Charity (or action) evaluation
Making recommendations of specific charities (or actions to take) to achieve a goal, such as the work done by GiveWell or CE.

Strategic decision-making (about causes):
Other cross-cutting research that is not directly on a specific cause but might prove useful to cause prioritization work. Examples could include historical analysis of who has done cause decision-making in the past and how it went, generally forecasting the future, comparisons of the evidence for different strategies for driving change (e.g. policy versus corporate campaigns versus values spreading) and so on. This could be similar to the supporting reports we have previously created at CE.

Section 3 maps out the promise of these different areas according to factors including neglectedness, impact potential, and expected difficulty. It explains in more depth our recommendation that a new organization should focus on strategic decision-making and cause candidate research.

2 Why Exploratory Altruism?

Key paths to impact

The impact on the EA community of finding new highly impactful causes could be huge. It could significantly increase the impact of the EA community and lead to large amounts of funds, time, and other resources being used in ways that are many times more effective than at present.

A new cause area could also broaden the EA community. A cause area that is particularly impactful under some worldviews or ethical views not yet dominant in the EA community could bring in new supporters who would otherwise bounce off EA. There is good precedent for this: EA research into animal charities was soon followed by the effective animal advocacy movement, and many animal advocates take EA ideas more seriously as a result of EA's attentiveness to animal issues.

Other advantages include:

- The creation of online resources. Cause reports could be combined into overview posts or handbooks like this Cause X Guide.
- Generally improving EA epistemics and ensuring that we are continuously checking our beliefs and assumptions.
- Possible cross-cutting advantages and mutually beneficial relationships if EAs are working on a broader selection of causes.
- Improved collaboration with actors outside the EA movement. EA can be somewhat insular, and could benefit from engaging with different causes, disciplines, and institutions.

This idea's impact is somewhat dependent on how likely you think there are to be other equally or more promising cause areas outside of the ones already identified by the EA movement.

Neglectedness

Judging by the rate of progress to date, cause candidate exploration in particular seems quite neglected. Currently new cause candidates seem to be explored by an individual EA community member or EA organizations bringing in a steelman for a single idea, such as economic growth, mental health, and preventing malevolent world leaders. However, this is not done systematically nor to a large number of potential areas. New cause areas are also occasionally brought in by outsiders to the movement who often lack the in-depth knowledge of EA prioritization necessary to most effectively evaluate the cause from an EA perspective. (This is true even for promising cause areas.)

Judged on other organizations' plans it is less clear how neglected this is. Organizations such as Rethink Priorities, Founders Pledge and Open Philanthropy all do some amount of exploratory research, including research into crucial considerations (e.g. sentience or invertebrate welfare) and possible causes (e.g. charter cities or nuclear security). However, none of these organizations have as their primary focus cause candidate research and our conversations with experts (including stakeholders at these organizations) suggest there is a gap for a new organization to enter the space.

Due to our concerns about replaceability, we reached out to Rethink Priorities (RP). They clarified that they expect there to be more than enough cause candidate work for another organization to conduct.

An organization solely focused on cause candidate research (and strategic decision-making) would ensure consistent rigor and descriptions which would support the cross-comparison of ideas.

Promise compared to other EA meta ideas

Within the space of possible EA-meta organizations, we think this specific idea is likely to be one of the most promising.

1) This research has a very high potential impact
A significant amount of money and time is moved as a result of EA cause research. A new cause or sub-cause could be many times more effective than current cause areas of focus. This research has the potential to shift resources in ways that are many times more effective. Find further details in our CEA, discussed below.

2) This helps address 3 of the 7 main problems identified in our community leaders interviews
In CE's interviews of forty EA community leaders, we asked about concerns with the EA movement that seemed promising for new organizations to tackle. Seven specific concerns came up 10 times or more. Exploratory research could help address 3 of these 7 main concerns: it directly addresses "closed-mindedness and intellectual

stagnation" and could also help to address "overconfidence in EA" and "reinventing the wheel".

3) Cause candidate research is the most promising EA-meta focus area based on our EA community leaders interviews

EA exploration was the most high ranking subarea to investigate. A lot of interviewees commented that this is core to EA as a whole and the framing of EA as a question. Many also felt this is quite neglected and done currently on a more ad hoc basis. Interviewees felt there was more room for systematic cause X research.

4) There is a clear need and interest in this research based on our interviews with relevant experts

Follow-up conversations identify a strong need for more exploratory research among EA organizations such as Open Philanthropy, 80000 Hours and Charity Entrepreneurship. We feel there are clear pathways for this research to progress along.

5) Cause research may be the most successful EA-meta intervention to date

The EA community has successfully conducted cause prioritization research in the past, so there is a precedent for this area being somewhat tractable. Depending on how exactly you divide up EA meta activities and how much you think different causes differ in effectiveness, it could be that cause prioritization research is the single most impactful thing the EA community has done to date.

6) There are good external precedents for exploratory research teams providing value and driving impact

In the public sector, many think tanks carry out research on the topics they think are most valuable for decision makers (primarily governments), and then promote this research to drive change. Many think tanks have an extremely broad cross-cutting remit and successfully run interdisciplinary research teams that cover a wide range of areas. In the private sector, consultancies provide supporting research to help decision makers choose where to focus resources, and also run research teams that can cover a broad set of topics.

7) Exploration builds on the research skills of community members

The EA community has a significant pool of potential research talent that it can draw on (based on recent hiring rounds of organizations such as Rethink Priorities and the Future of Humanity Institute).

8) Our EA-meta crucial considerations analysis suggests this idea is reasonable.

In our interviews with EA community leaders we considered the following crucial considerations.

- Improve rather than expand the community. 42% of interviewees thought it was better to improve, 35% to expand, and 23% were unsure. This project focuses on improving, although new causes may expand the community too.
- Focus on money rather than time or information. 34% of people thought it was better to focus on money, 26% on ideas, 23% on talent, and 17% were unsure. This project focuses on ideas.
- Broad rather than narrow EA-meta organizations. 41% leaned towards broad organizations, 32% towards narrow and 26% were unsure. This organization is broad and may even expand the scope of the EA community.

This project matches 2 out of 3 of the considered crucial considerations.

9) This idea is sufficiently flexible to the changing needs of the community
High-level strategic research is useful for EAs across different worldviews and a skilled interdisciplinary research team should be able to move between topics as needed. It is also an organization that could do equally well or better with large scale changes of the EA movement.

Our main concerns

1) Execution difficulty
Since the research has a very broad scope and is relatively challenging, it would be difficult for founders to do well. Staff would need a very broad understanding of a vast range of different topics. Prioritizing, strategizing, and justifying decisions to donors would be very important. This is discussed more in Section 3. A. Research approach, below.

2) Counterfactual replaceability
It is a truth universally acknowledged, that an EA in possession of a good fortune, must be in want of a cause. As such, cause candidate research might be conducted even if this new organization were not founded. Currently, our research suggests that there is space and need for a new organization doing this work and in general there is room for multiple organizations doing this sort of work. However, there is a risk that founding a new organization in this space could take talent away from or otherwise compete with adjacent organizations (e.g. RP and FHI), weakening all

three organizations. We note that RP have advised that they support this idea and that there is a lot of room for another research organization like RP to exist.

3) Hidden costs and risks of research
Research showing that the case for a popular cause is weak could cause controversy within the movement. Research showing that the case for a new cause is strong could shake up the movement and increase levels of disagreement between EA actors. Low quality or misleading research could prove a distraction to the broader EA community, and could shift funds to a new cause area that is less effective than known cause areas. This is discussed more in the section on Managing risks, below.

Additional notes

CEA

Our cost-effectiveness analysis estimates the increased impact of funds donated based on this organization's research. Impact is monetized in the form of equivalent donations to today’s current top-tier EA charities. Our estimates also suggest that the impact of this organization will increase over time. For example, in year 3 the estimated impact is as follows:

Increased impact of donations (monetized)	$ 1,500,000
Costs	$ 200,000
Leverage ratio	~8

Please note this CEA is only comparable to other CEAs created for EA meta charities by Charity Entrepreneurship. Other EA organizations use vastly different methodologies to estimate their leverage ratio.

Expert interviews

In our in-depth interviews, experts were generally positive toward this charity idea. Many experts agreed that more exploration of possible cause areas from an EA lens could lead to new results. Many also expressed concern at how hard it was for the movement to consider a new cause relative to the established ones.

Research potential
Experts felt that it was fairly likely that the EA movement would focus on different

and adjacent causes in the next 10 years. They estimated between 1 and 30% chance that the EA movement is focused on very different topics in 10 years, and 35 to 90% chance that the EA movement is focused on adjacent topics in 10 years. There are still many causes in both neartermism and longtermism (easily at least twenty in each) that could be investigated.

Appetite to use exploratory research
Experts felt that Cause Candidate research was quite promising for both neartermist and longtermist cause candidates. For neartermist causes, large donors are open to new neartermist areas. When they enter a new cause, they commit to that cause for several years and a significant financial amount. We think they, and other large neartermist donors could be receptive to research. For longtermist causes, there are some meta/longtermist organizations that could be consumers of cause candidate research and may promote other causes to the community if they are convinced by good arguments.

Key areas of research

- Most experts who commented on neartermism felt that EA has only covered a very small portion of potential interventions. Experts also mentioned researching neglected geographic regions (from a neartermist perspective) could be very important to help find potential blind spots.
- Experts mentioned improving the ITN framework would be useful as it has been shown to have flaws. One mentioned that neartermist causes have been incorrectly excluded because of it.
- Policy was cited by many experts as a promising focus area.

Key uncertainties
The main hesitations and uncertainties about this project were related to the specific things this org could research. For example, one expert said that the biggest challenge would be to sufficiently narrow the scope of the org. Experts thought that worldview research in particular could end up being very difficult and intractable, or could be done wrong. Another concern was whether the organization could find a niche given that other organizations are working on similar projects. These uncertainties influenced our recommendation – our research suggests that cause candidate exploration is particularly narrow, tractable, and neglected.

3 How to run an Exploratory Altruism organization

The scope of exploratory altruism could be incredibly broad; we strongly recommend that the organization pursue a narrower approach to fill the most important gaps. At a very high level, the theory of change for this organization will be:

A. Do research

+

B. Ensure research is impactful

↓

Impact

The sections below map out the two broad steps of the theory of change. We first map out the 5 different categories of research. Then we map out the 3 potential audiences and 2 influencing models to ensure impact. The strategy implemented and audience targeted will likely depend on the co-founders and results of experimentation. For each category, we have evaluated the options using an informed consideration methodology. Summary tables are provided in the relevant sections.

A. Research approach

Types of research – mapping

The table overleaf maps out the different types of exploratory research (introduced in Section 1) that a new organization could do:

	Worldview Research	Strategic Decision Making about causes	Cause Candidate Research	Intervention Research	Charity Evaluation
	Broad research into what matters most in the world	*Investigating different approaches for doing good*	*Comparing causes within worldviews*	*Comparing interventions within causes*	*Comparing charities with-in causes or interventions*
Ideal Target Audience	EA	EA and non-EA	Mostly EA	EA and non-EA	EA and non-EA
Strength of Past Successes	Somewhat good (e.g. GPI and Open Phil)	Somewhat weak (e.g. FHI, individual researchers, GJP)	Quite good (e.g. RP, Open Phil)	Very good (e.g. RP, FP, CE, GiveWell)	Very good (e.g. ACE, GiveWell, FP)
Current Neglectedness	Very neglected	Somewhat neglected	Very neglected	Somewhat neglected	Somewhat neglected
Expected future neglectedness	Somewhat neglected	Somewhat neglected	Somewhat neglected	Somewhat neglected	Very neglected (as more causes identified)
Potential impact	Very large	Very large	Large	Large	Very large
Feedback loop	Somewhat weak	Somewhat good	Quite good	Good	Good
Co-founder skills	Very important	Very important	Moderately important	Somewhat important	Somewhat important
Interest from EA orgs	Moderate	Moderate-high	Very high	Very high	Moderate
Information value	Very high	Very high	High	Low	Low
Expected difficulty	High	High-moderate	High-moderate	Moderate	Moderate

Columns in grey are research areas we do not recommend that this organization pursue.
(Acronyms: ACE – Animal Charity Evaluators, CE – Charity Entrepreneurship, FP – Founders Pledge, FHI – Future of Humanity Institute, GPI – Global Priorities Institute, Open Phil – Open Philanthropy, RP – Rethink Priorities, GJP - Good Judgment Project)

Our recommendation

We would like to see a new organization primarily focusing on cause candidate research (e.g. how does economic development compare to global health as a cause area aiming to address human suffering/poverty?). We also envision the new organization conducting some degree of strategic decision-making research to explore new areas within existing causes (e.g. how promising is policy compared to direct charity work compared to impact investing?). Based on our expert interviews there are clear EA organizations who would use or be interested in both these kinds of research.

We do not recommend charity/intervention research because it requires a greater degree of specialization, although we would expect some charity and intervention research to be necessary as new cause areas are identified. When assessing a cause area, it would be useful to identify promising interventions and charities and to make practical recommendations to users of the research as to how they could invest resources into supporting or building up this cause area.[1]

Although we think worldview research is very valuable, we do not recommend it because it is highly difficult & uncertain. That said, we expect that some of the research would likely be useful for worldview level decision-making.

Choosing what specific topics to research should also depend on the expected path to impact and audience for the research, as set out in Section 3.B. Ensuring research is impactful, below.

Overall cause candidate research is certainly a broad enough area to warrant an organization exclusively focused on this process.

Cause candidate research

Cause evaluations

We expect the simplest way to do cause candidate research would be to choose a single cause at a time and to carry out an in-depth evaluation. The organization

[1] The new organization could perhaps partner with Charity Entrepreneurship for such research. This could create valuable synergies and ensure that the new organization's research is translated into real-world impact through the founding of new charities.

would then present the case for believing that cause is the most important cause for a significant amount of altruistic people to be working on.

The evaluation could contain:

- a steelman of the case for the cause
- an analysis of the reasons for or against the cause
- a synthesis of different kinds of evidence about the value of this cause – expert consensus, historical traction, scale/neglectedness/tractability analysis, philosophical or economic analysis, case studies, etc.
- recommendations for how to progress the cause – key interventions, charities, or policies that look promising

The organization could evaluate cause areas adjacent to EA which are seen as promising by some but not the majority of EAs, such as the Problem areas beyond 80,000 Hours' current priorities by 80,000 Hours for longtermist causes or the list of causes Open Philanthropy is investigating in neartermism. These areas are often talked about already but lack detailed research comparing them to their closest reference group (e.g. Global health - Mental health / extreme pain / economic development, Farmed animals - wild animals, Environmentalism - Global poverty - Farmed animals, Improving institutional decision-making - Global peace / cooperation / Reducing existential risk). These might be the causes most likely to gain major support in EA. However a much larger amount of research would need to be done to deepen these areas as many of them have entire orgs already focused on them. We estimate it would likely take 3+ months per area to make meaningful extra contributions to each of these areas. ~18 months worth of research

The organization could evaluate cause areas that might be promising but so far have been relatively neglected by the EA movement. These ideas would be more novel and less likely for the EA movement to consider absent the meta org e.g. trade liberalization, reducing loneliness, meta science and epistemology. This is a broader list of cause candidates here. These cause areas could have major progress made on them relatively quickly however likely will have a lower hit rate than the more established areas. There are at least 20-40 ideas that could be evaluated. After the first round of shallow research, the list of detailed reviews could be between 5-10. Meaningful progress could likely be made in 2-3 months per cause area. Some cause areas have considerable momentum outside of the EA movement and thus if made more effective could lead to larger scale changes. For example, education reform, disaster relief and human rights violations. These areas will likely carry a lot of outreach potential if they are found to be impactful. From a movement building perspective, such research can be useful to share with newcomers who may be very interested in their specific cause areas.

Cause comparisons

An alternative, and potentially more complicated way to do cause candidate research would be to compare between causes. Once a number of cause explorations are complete the organization could pick two or more causes and find ways to compare between them.

This might require choosing a particular worldview to support comparison. Yet this could be done in an exploratory: by choosing worldviews that have had less of a focus within the EA community to date (e.g. prioritarianism, justice-focused ethics, non-x-risk focused longtermism) and exploring what causes people with that worldview might gravitate towards. This kind of work crosses over into strategic decision-making.

Strategic decision-making about causes

We are less certain of the promise of strategic decision-making than that of cause candidate research. However, we think this could nonetheless be a promising avenue for a new organization and that cause candidate research could benefit from and be supplemented by analysis of different approaches for doing good. This could look like:

- Decision making – e.g. how to make decisions given uncertainty? How important is it for a path to doing good to have clear feedback loops?
- Analysis of other organizations/movements – e.g. historically have groups who tried to maximise their impact been successful? How do other organizations make altruistic decisions?
- Strategies – How do different strategies for changing the world affect cause decisions? How promising is policy compared to direct charity work compared to impact investing? How does this change under different worldviews?
- Supportive research – e.g. research and apply futures methodologies (such as forecasting).

Research into how to pick causes using different strategies could be valuable for reaching EA ideas to new audiences. Traditionally most EA cause research has been focused on donating which can lead to people with other advantages, such as skills in policy making or campaigning, bouncing off the movement.

Execution difficulty – Jack of all trades, master of none

We think execution difficulty is a challenge that the founders will need to consider carefully. Ideally the research team should have a mixed set of skills and interests. Having a strong and consistent process would be key for creating cross-comparable outcomes, although this needs to be balanced against ongoing learning and having a process that continually updates over time. Further research about how best to approach this challenge is needed; there are a few interdisciplinary research teams within the EA community who may be able to provide advice.

B. Ensuring research is impactful

Doing the research is only half the battle. To demonstrate impact an exploratory altruism organization needs to ensure that the research is used by decision makers.

In this section we explore three possible theories of change:

i.	ii.	iii.
Do research + *Make it public* ↓ *Influence EA readers*	*Do research* + *Build a new non-EA audience* ↓ *Influence new audience*	*Do research* + *Persuade other EA orgs to use your research* ↓ *Influence their audiences and money*

i. Influence EA readers

Why this is promising

The EA community provides a reasonably receptive audience for cause prioritization research. Past cause analyses on topics such as mental health and economic growth have been picked up by various EA community members. This has led to donations to new charities and new EA sub-communities focusing on or working in these areas.

How to do this

This could simply involve posting work on the EA Forum. In some cases the Exploratory Altruism founders might want to support sub-communities being formed, or support individuals looking to do further research on interventions.

Tracking impact

Donations to suggested charities could be tracked through affiliate links or by talking to the charities and funders. New EA sub-communities could be identified. The percent of EAs who mention each cause on the EA survey could be tracked.

ii. Building a non-EA audience

Why this is promising

As has happened with the animal and EA movements converging and creating the effective animal advocacy movement, we think an organization in this space has significant potential to reach new non-EA audiences. Based on our analysis we think this could be as valuable as reaching existing EA audiences.

	EA audience	Non-EA audience
Neglectedness	Neglected (only Open Phil and RP are doing some systematic cause prioritization research)	Very neglected (GPI targets non-EA academics, other orgs push specific causes)
Past success of this approach	Moderate. Rethink Priorities has built traction with Open Phil. This success is not counted in our CEA as not yet included in RP impact reports	High. GiveWell, Founders' Pledge, ACE
Potential impact	Very high. Could result in significant resources invested in a new cause area if Open Phil commits to it	Very high, but with high variance - depends on which cause is chosen

How to do this

It would be problematic if the Exploratory Altruism organization pivoted and focused solely on one promising cause area after identifying a promising area, as that would leave a gap for exploratory research which is crucial. However, smaller steps could be taken to grow a cause area beyond EA.

There is a strong precedent for new charity evaluators in new causes to attract significant attention and funds. For example, Animal Charity Evaluators preceded and helped seed the effective animal advocacy movement. An Exploratory Altruism organization could:

- Set up funds, with committees made of experts, in the identified new cause area (such as the EA Funds or the Founders Pledge funds). We expect this would be significantly less work for the organization than starting a charity evaluator.
- Support EAs and/or people working in the identified cause area to start or spin off a charity evaluator or an intervention research organisation in that area.

There is also some precedent for EA work being picked up in other domains. For example, Charity Entrepreneurship's research on animal welfare has been cited in a legal case in the US. An Exploratory Altruism organization could promote research to relevant decision makers, e.g. through partnering with think tanks or policy advocates.

Tracking impact

Money moved would be a key metric, and could be tracked through affiliate links, by talking to the charities and funders, or by moving money through the organization. Reach and readership of ideas, or policy changes could also be tracked.

iii. Influencing existing EA organizations (and securing funding)

Why this is promising

As we set out in the Expert interviews section, experts from both longtermist and neartermist EA organizations have expressed a significant interest in more exploratory research. They are capacity constrained internally and keen to use high quality research from other organizations should it present itself. Furthermore, producing research for large EA organizations could also be a path to securing funding for an fledgling Exploratory Altruism organization.

How to do this

We have broken this down into two different organizational models:

- **The consultancy style approach**. An established EA organization identifies challenges that it has that research could solve, vets the Exploratory Altruism organization, and funds them to do the research.
- **The think tank style approach.** The Exploratory Altruism organization (funded by a donor) identifies challenges at established EA organizations and

produces research targeted at addressing those challenges and persuading the established organization to adopt the suggested solutions.

	Consultancy style model (stakeholder driven research)	Think tank style model (expected-value driven research)
Examples	RP (about 50% of their projects are for Open Phil)	GPI (influences academia), RP (about 50% of their projects)
Flexibility	Moderately low (may be restricted to specific topics)	Very high
Potential impact	Very high	Very high
Chance of success	High	Medium
Chance of funding	High	Medium
Reliance on co-founders' skills	Moderate (research is likely to be more concrete and less abstract)	Very high (need skills to choose the right topics & funders need to trust them)
Counterfactual replaceability	Moderate (RP is doing this, and scaling quite quickly)	Very low (there are gaps in current research)

Overall we tentatively recommend the think tank style model, partly because Rethink Priorities appears to be leaning towards the consultancy style model. Additionally, outside of EA, the think tank style model is the dominant way of improving the knowledge and capability of organizations that are trying to do good. Partly we think a think tank model will come to more novel ideas than a consultancy based approach.

In both cases, it would be important for the co-founders to have strong research backgrounds. For a think tank, research direction would be more informed by the co-founders' judgments of what is important. Even with this model, close engagement with established organizations would be necessary to ensure the research is targeted to where it can have an impact and is practically useful.

Tracking impact

Assuming there is a decent relationship with key EA stakeholders, the use of research by those stakeholders, and resultant changes in decisions or money moved could be tracked.

The optimal theory of change

The optimal theory of change would depend on the co-founder's background and profile. The actual strategy implemented will likely be a combination of the different audiences and strategies listed above. The strategy may change as the organization builds its brand and reputation. The founders will need to experiment, iterate, and adapt to find the ideal strategy.

The target audience might also depend a lot on the specific cause area or research topic under consideration. As the organization moves between topics it may move between audiences. Some topics may be very relevant to existing EA organizations and some of most interest to actors outside the EA community. The founders would need to evaluate the expected difficulty of outreach to each target audience and determine where to invest resources in trying to influence different audiences.

Founders will need to have sufficiently granular feedback loops in place to track the impact of various pieces of research, and learn and adapt their approach accordingly.

Managing risks

- Maintain robust quality assurance to ensure research is high quality
- Test the way research is interpreted and understood by the community to avoid misleading or confusing the community as a whole
- Ensure that research is contextualized correctly, so that readers can compare it to past research on related topics and update their views accordingly (i.e. try to make apples to apples comparisons where possible)
- Take a charitable view of past cause considerations

Other practical advice

- Have a focused agenda that is legible to funders and stakeholders
- Get good at understanding the model of thinking of different stakeholder organizations
- Understand why other organizations have chosen different topics
- Explore funding opportunities outside of EA, especially if researching topics that are popular in the mainstream

Founder personal fit

Co-founders would need a strong and well-rounded background in research, ideally with complementary interests. This would also make it more difficult to influence stakeholders with the research. Solid communications skills would also be necessary to identify and influence stakeholders. A final criteria would be a genuine openness to finding cause areas outside of the current EA top ones without specific loyalty to a single idea.

As the organization grows, another challenge would be to find other researchers with experience in a wide range of topics (although CE and RP have managed to do this in the past).

4 Next steps

We recommend an exploratory altruism charity as one of the meta charities we would be keen to incubate. We believe this idea can have a positive influence on the trajectory of the EA community and direct funds to effective activities. Next steps for this organization would likely entail more closely considering the different paths to impact and the skills of specific co-founders.

Appendix A: Further resources

On GPR research / Cause X research

- Prospecting for Gold
- Draft: The Case for Cause Prioritization
- Why global priorities research is even more important than I thought
- Is region-level cause prioritization research valuable to spot promising long-term priority causes worldwide? - EA Forum
- Objections to Value-Alignment between Effective Altruists - EA Forum
- Tips on starting new org from Michelle Hutchinson
- Cause candidates - EA Forum
- The case of the missing cause prioritization research.
    - Michelle Hutchinson's comment

Appendix B: Acronyms and abbreviations used

ACE – Animal Charity Evaluators
CE – Charity Entrepreneurship
EA – Effective Altruism
FP – Founders Pledge
FHI – Future of Humanity Institute
GPI – Global Priorities Institute
GW – GiveWell
Open Phil – Open Philanthropy
Org – organization
RP – Rethink Priorities

Annex: Research methodology

Research to determine our EA meta charity recommendations

Many weak arguments

In spaces with less evidence available, we tend to favor many weak arguments over a single strong one. As such, we look at how to improve the EA movement from many different angles. Some of these fall within traditional EA frameworks, such as cost-effectiveness analysis (CEA); others are less commonly used in EA but still apply, e.g. limiting factor analysis (which identifies bottlenecks that will impede progress on an issue), or theory of change models. We can expect less consensus about the most important interventions in EA meta than in more established cause areas, although we can still define tiers of weaker and stronger ideas.

Prior views

These include the views of our board, advisors, and CE staff members on how to improve the EA movement specifically and social movements in general.

Our team has been deeply connected to the EA movement for many years. We have gained insight into how to improve the movement through projects we have worked on or consulted with outside of our formal research years (such as helping advise meta projects in the EA space) as well as through direct research prior to founding Charity Entrepreneurship.

More recently, meta research featured in Charity Entrepreneurship's animal advocacy work. Our recommendation of Animal Advocacy Careers improves animal welfare indirectly, through strengthening the animal advocacy movement. This past meta research on animal advocacy informed our approach to our EA meta research.

Synthesized views

To incorporate the varying perspectives from across the EA community, we surveyed a number of different EAs including chapter leaders, meta EA funders, and individuals working full time at EA meta organizations. The survey asked about broad crucial considerations and about specific areas or ideas that might be promising. This gave us a soft sense of the ranking of a wide range of ideas, which we narrowed down using our traditional method of iterative depth.

In addition to this survey created in-house, we drew significantly on the yearly EA surveys (conducted by Rethink Priorities) to better understand trends and gaps. These sources are the closest thing to hard empirical data that the EA movement has on itself. We also used some isolated data sets that more specifically targeted a key question (e.g. on value drift).

Cross applicable data

We pulled out information from other cause areas that better track and evaluate their own impact, including the animal and global health movements. In studying these other movements we looked at solutions as well as mistakes and how to avoid them. Although we think the EA movement is unique in many ways, it still seems likely that information can be cross applied, particularly where common failure modes exist.

Crucial considerations

A large part of our research into EA meta interventions looked at cross-cutting crucial considerations, to rule out or highlight as promising multiple ideas. For instance, a consideration such as how important it is to grow versus improve the EA movement could greatly affect the prioritization of different ideas. We have pulled these considerations into a separate report. Key considerations include:

- Is it better to expand or improve the EA movement?
- Does the EA movement need more time, money, or information?
- Should the EA movement be broader or more narrow?

Flexibility

We considered and weighted highly two types of flexibility: organizational and movement flexibility. In a space like EA where evidence is scant and the movement is rapidly changing, it seems important to create organizations that can achieve impact even as new trends or perspectives arise. This tends to result in flexible organizations whose focus is slightly broader or whose approach can be easily adapted to changing circumstances. Organizational flexibility leads to movement flexibility – the movement as a whole can more readily grow and improve over time. For example, some meta organizations within EA would do equally well if a new cause area were added, while others would fare worse. We think that this can create negative norms and intellectual stagnation and thus put a value on organizations that can adapt easily to this sort of update. We expect having organizations like this will tend to help the EA movement improve over time.

Additional research for developing this charity recommendation

After identifying three EA meta charities to recommend, we carried out further in-depth research to develop the ideas. This included

- Informed consideration. Broad research thinking through each idea from many different angles, brainstorming potential approaches, and considering crucial considerations.
- In-depth expert interviews. We spoke to three to five experts who might have good views on the idea and could provide a deeper dive than our initial interviews. We sought out people in the EA space who would have views on exactly what the need and appetite is for a new charity and on what the charity should focus on.
- Theories of change. We mapped out a few theories of change for each of the charity ideas, comparing them and drawing conclusions about the possible ways the new charities could be run.
- Cost-effectiveness analysis. We carried out a very rough cost-effectiveness analysis for the plausible impacts the charity could have for few of the most promising theories of change that could be adopted.

We spent between 2-10 hours on each of these steps. This additional in-depth research on each idea was combined with our earlier cross-cutting research and written up into the report you see here. The report does not include private notes, and we will speak to the founder of this charity in far more depth than is provided in these pages.