Evaluation report — Longview's Emerging Challenges Fund

Written by: Michael Townsend
Published: November 2023

As part of Giving What We Can's (GWWC) work evaluating evaluators, this report lays out our evaluation of Longview Philanthropy's Emerging Challenges Fund (previously, the "Longtermism Fund") as a grantmaker in reducing global catastrophic risks (GCRs), and explains why we decided to recommend the ECF as a top-rated fund and plan to allocate half of our Risks and Resilience Fund's budget to the ECF as of November 2023.

For our evaluation, we reviewed the Emerging Challenges Fund's (ECF's) and Longview's (LV) public communications, and were provided access to some private materials, such as the reasoning behind a selection of grants, and various other internal documents.

In our assessment of the ECF, we also looked at LV's grantmaking in general. This was largely motivated by our expectation that some of the effect of donations to the ECF will be to free up resources for LV's other grantmaking. We additionally tried to compare the ECF to the other evaluator we evaluated that works on reducing GCRs: EA Funds' Long-Term Future Fund (LTFF).

There were several limitations to the evaluation. One we would like to highlight upfront is our conflict of interest: we had initiated the fund in 2022 and asked LV to manage its grantmaking, and we had since managed its communications, writing the copy for the launch of the fund and its subsequent grants reports. Given our new research direction in evaluating evaluators, we were motivated to reduce this conflict. We now no longer manage these communications, and the fund is completely owned by LV and has since been renamed.[1] Other limitations include that we had limited access to relevant information (in part because the ECF and LV were not set up to be evaluated in this way), and we were highly time constrained.

The key findings informing our recommendation of the ECF are:

- Longview has solid grantmaking processes in place to find highly cost-effective funding opportunities.
- In the grants we evaluated, we generally saw these processes working as intended, which makes us optimistic about the cost-effectiveness of the grants.
- The scope and structure of the ECF is — by design — consistent with what we are looking for with our Risks and Resilience Fund: a fund that makes grants that are

[1] This transition and renaming will be implemented on our website and we sent an update to all previous donors of the Emerging Challenges Fund on November 28th. To be clear, this does not represent a material change to the fund's grantmaking. It mainly means that Longview will from now on publish its own grant reports, rather than Giving What We Can publishing reports with help from Longview.

relevant and understandable to a wide variety of donors looking to reduce global catastrophic risks.
- We don't know of any clearly better alternative donation option in reducing GCRs.[2]

We found no clear, justifiable reasons for the donor's extra dollar to be better spent at the ECF than at the LTFF (or vice versa). As a result, we recommend both as top-rated funds and plan to allocate half of our Risks and Resilience Fund's budget to each until our next evaluation. We did outline several differences between the ECF and the LTFF so motivated donors can decide for themselves which they think fits their values and starting assumptions best.

This report contains:
- Our framework and approach
- Our takeaways from the evaluation
- How we think the ECF compares with the LTFF
- A conclusion summarising our takeaways and laying out our plans for a future evaluation

Our approach

This section provides more detailed context on the approach and framework we used for the evaluation, including why we chose to investigate the ECF and the limitations of our evaluation.

Why are we evaluating the evaluators?

We want to help donors maximise their impact by recommending the best funding opportunities from a range of (sufficiently plausible) worldviews. To do so, we work with evaluators we selected as being best able to help donors in this way,[3] and so the goals of this project are to:
- **Understand** the strengths and weaknesses of different evaluators.
- Use that understanding to **decide** whether and how we defer to the evaluator's recommendations (ideally in direct comparison with at least one other high-potential candidate in the same cause area).
- Transparently **explain and justify** that decision to donors.

More concretely, this project informs:
- Which evaluators we rely on for our charity and fund recommendations.

---

[2] Note that we haven't yet evaluated Founders Pledge's Global Catastrophic Risk Fund, but aim to do so next year.
[3] At the beginning of our evaluating the evaluators project, we considered the relevant bar as to whether the evaluator "reliably recommended and/or granted to the most cost-effective funding opportunities based on a sufficiently plausible worldview." We received pushback from external reviewers that, taken literally, this was an overly onerous bar that no evaluator could claim to meet. The issue was that we were using "most cost-effective" in a confusing way. We intended it to mean (roughly): "compared to other available recommendations or funding opportunities that we have evaluated, there is no clear superior alternative that seems more cost-effective, ex-ante" Given this is a bit of a mouthful, we decided to explain our approach as we've done here: selecting evaluators that we think are best able to help donors maximise their impact.

- Which evaluators we ask to advise our cause area funds.

Just as we did with our own impact evaluation, we used usefulness, transparency, and justifiability as our guiding principles (as opposed to comprehensiveness and procedural uniformity). In this vein, we didn't attempt to come up with an overarching framework or model for evaluating the evaluators — instead, we focused on trying to answer the questions that seemed most decision-relevant for reaching our goals of understanding the evaluator, deciding whether to defer to their recommendations, and being able to explain and justify that decision. (You can see our list of starting questions in the appendix.)

We go further into why and how we evaluate evaluators on our website.

Why did we want to evaluate the ECF?

The ECF had previously been one of our recommended funds, and we had been involved in its communications and inception (see below for more on how we dealt with this conflict of interest). At the time, we initiated the fund because we thought there was a need for a funding option for donors looking to support larger, more easily understandable projects working to make the future go well. We saw this fund as a nice complement to the LTFF, which also supports smaller projects which are sometimes harder to understand without specific context on the cause area. We had asked LV to manage grantmaking for the fund based on its reputation in the effective giving ecosystem and its alignment with the goals of the fund we envisioned.

We thought it would be worth investigating the ECF because:

- We wanted to form a richer understanding of LV's work and check whether it (still) aligns with our purpose of helping donors maximise their impact in reducing GCRs.
- We conducted a survey of fundraising organisations to help us decide which evaluators to prioritise investigating. We asked which evaluator they **don't** currently defer to that they would consider deferring to if we reviewed them and came to positive conclusions. Six of the 16 organisations named LV (second only to Happier Lives Institute).
- We thought the ECF was a contender to be the most cost-effective funding opportunity for donors looking to reduce GCRs, but we weren't sure how it would compare to the LTFF on further analysis. We thought it could be helpful for us to compare the two — even if we could not determine which one was a more cost-effective option, we thought we could provide considerations that may help donors choose between them.
- The ECF was set up specifically to be a default giving option for a wide range of donors, so it seemed like an especially strong fit for our Risks and Resilience Fund.
- We thought our evaluation could be useful to the ECF itself — and to other evaluators and grantmakers working to reduce GCRs — by highlighting what we think the ECF does well and where we see room for improvement.

Limitations to our evaluation of the ECF

There were various limitations to our evaluation of the ECF. In this section, we discuss these and how we have tried to address them.

The first of these was our conflict of interest. We initiated the ECF in 2022, asked LV to manage its grantmaking as a trusted evaluator, and were responsible for the fund's communications. To reduce this conflict of interest, we have stepped down from managing the fund's communications, and the fund now operates independently of GWWC and will continue under a soon-to-be-announced new name.[4] This decision was made partway through the evaluation.[5] Another conflict is that we are both projects of the same legal entity, Effective Ventures.

Second, the ECF was not set up to be externally evaluated in this way:

- We were only able to access a limited number of its past grants.
- Even for some of the grants we had access to, key considerations were anonymised due to them containing sensitive information.

Third, our research team has limited expertise in grantmaking and evaluation to reduce GCRs. Though both Michael Townsend (the primary author) and Sjir Hoeijmakers (who provided extensive input on this report) have some direct experience with grantmaking and evaluation to reduce GCRs, we don't have deep subject-matter expertise across this space. Given we needed to make many subjective judgement calls throughout this report, this was a significant limitation. We tried to account for this to some extent by asking external reviewers for input, though in future we would like to seek input from experts earlier in the process, rather than just at the review stage.

Fourth, the evaluation and grantmaking space in reducing GCRs is not especially developed (particularly compared to global health):

- Experts in the field often disagree on whether a grant is positive or negative.
- It is hard to find direct feedback loops on which interventions have had which effects.
- Much of the work being funded is relatively new, and therefore there is often a limited track record to refer to and no widely agreed-upon metrics.

Fifth, we were highly time-constrained in completing this first iteration of our evaluating the evaluators project. On average, we were able to spend about 10 researcher-days per organisation we evaluated (i.e., two researchers spending one full workweek), of which only a limited part could go into direct evaluation work — a lot of time goes into planning and scoping, internal discussion on the project as a whole, communications with evaluators, and so on.

---

[4] This transition and renaming will be implemented on our website and we will send an update to all previous donors of the Emerging Challenges Fund around November 27 (we will also edit this report to reflect the new name). To be clear, this does not represent a material change to the fund's grantmaking. It mainly means that Longview will from now on publish its own grant reports, rather than Giving What We Can publishing reports with help from Longview.

[5] At the outset, we were not clear whether we would be evaluating LV **generally** (i.e., all of its grantmaking — not just for the ECF specifically — and also whether it could provide charity recommendations), just LV as a grantmaker, or just the ECF. In part, this was because we weren't sure about the future of the ECF at the time (we were considering various options for the fund, which would have implications for our evaluations). Once LV agreed to taking full ownership of the fund, we were in turn slow to clarify the scope of our investigation, which was a mistake we'd like to acknowledge.

Given these limitations, we aimed to:

- Focus our efforts where they would be most **useful**, by prioritising decision-relevant considerations over those without immediate relevance to our recommendation decisions.
- Only make decisions we thought we could **justify**, and highlight our uncertainties where we didn't feel confident enough to make a decision.
- Be as **transparent** as possible in our reasoning in this report, without breaching any confidentiality agreements we made and limited by our time constraints.

Even with our efforts to take a transparency, justifiability, and usefulness approach, we appreciate there still are significant limitations to our evaluation of the ECF, and see this as a minimum-viable-product attempt at such an evaluation which we look forward to improving on in future iterations. For instance, we have aimed to be transparent in our reasoning, but have had to trade this off against confidentiality concerns (e.g., we're not able to publicly share the details of individual grants) and usefulness (e.g., we focus on broader patterns rather than going deep into individual examples). This means that the reader will still have to trust our judgement and reasoning to a significant extent. We are not sure if we have struck the right balance here, and invite feedback from readers on this and other tradeoffs and choices we've made.

How we evaluated the ECF

We wanted to find out:

- Whether we should recommend the ECF to donors who want to maximise their impact in reducing GCRs (compared to other plausible alternatives like the LTFF).
- What LV is currently doing well and where it could improve.

To do this, we were interested in answering some of our general starting questions, shared in the appendix. In the case of the ECF, we were especially interested in:

- The quality of LV's grantmaking process
- The cost-effectiveness of past grants from the ECF
- How ECF compares to the LTFF

To help answer these questions, we looked at both public and private information.

Public information included:

1. LV's website
2. The ECF's communications

Private information included:

1. Partial access to some private information on all of the ECF's grants
    a. For several of these grants, we had (almost full) access to the case behind the grants.
    b. For the rest of them, we were able to see LV's communications to donors about the case for the grant. These documents generally contained fewer considerations than were internally considered.

2. Partial to full access to the reasoning behind six of LV's past grants or grant investigations, including additional documents of LV's followup to these grants for several of these.
3. LV's internal grants processes, including:
    a. Various questionnaires it provides to grantees.
    b. Its internal grants process document.
    c. An annotated list of its network and advisors in the focus areas it supports.
    d. Various other internal documents and information, including high-level descriptions of its grantmaking practices and the kinds of grants it has made.

We're grateful that LV was able to share this information with us and to invest significant time in answering our questions, and to do so while being respectful of the privacy of grantees.

We go further into why and how we evaluate evaluators more generally on our website.

Our takeaways from evaluating the ECF

In this section, we share our takeaways from evaluating the ECF. We start with the key reasons informing our recommendation of the ECF to donors, before we highlight some other findings that may be relevant to donors, LV, and other grantmakers.

Key findings informing our recommendation of the ECF to donors

We want to recommend the most cost-effective funding opportunities we can to donors. In the case of the ECF, it is insufficient to assess this purely by looking at the fund's grantmaking track record — partly because the fund is relatively new, but mostly because the fund is quite explicit that many of its grants will be highly fungible. That is, the fund often grants to projects that would have otherwise been funded (for example, by one of the philanthropists LV advises, or by another funder entirely). Given this, it is also important to check:

- Longview's grantmaking outside of the ECF.
- How well LV coordinates with other funders.

With this in mind, the reasons we recommend the ECF are:

- We think LV's grantmaking processes are generally aligned with LV making highly cost-effective grants in reducing GCRs (including by coordinating with other funders).
- LV's resulting grants seem plausibly highly cost-effective (though we're not sure how the ECF's cost-effectiveness compares to the LTFF specifically, which we discuss in the next section).
- The ECF's structure and scope is relevant and understandable to a wide variety of donors looking to reduce GCRs.
- We are not aware of any clearly better alternative donation options to reduce GCRs.

Below, we provide more detail on each of these points.

Longview's grantmaking process

We think LV's grantmaking process allows it to find, fund, and support highly cost-effective organisations reducing GCRs. Longview describes its principles and process at a high level on its website, though our views are also informed by internal documentation of this process and seeing the process used in practice in its grantmaking.

Longview's principles and goals

We think LV's principles — radical impartiality, intellectual honesty, hits-based giving, win-win scenarios, and a scientific mindset — and overall institutional aims align well for donors looking to maximise cost-effectiveness.

Access to highly effective funding opportunities

We were surprised by how expansive LV's network was, and how proactive LV was as a grantmaker in using that network to find and fund the most cost-effective opportunities. This is based on seeing:

- Its internal list of advisors and experts.
- A list of grants to projects where it played a significant role in the development of the project (e.g., by providing seed funding, or helping shape the scope of the project).
- A specific example of LV being highly involved in funding the creation of a new organisation.

In addition to this networking giving LV a wide range of potentially promising funding opportunities, we think that LV's filtering of funding opportunities reflects what we see as best practice in grantmaking. This involves starting with a longlist of potential options that are already determined to be promising in various ways — like working in an impactful cause area, or being recommended by their network — and filtering these down into only the *most* cost-effective. A high-level overview of this process is described on its website. In the grants we saw, we thought the process looked like it was functioning as intended, though we have limited information on how LV *excludes* grants in its process, as we only saw one case of a grant being thoroughly considered and then going unfunded.

Taking into account a wide range of important considerations

We think one of the key strengths of LV's process is that it takes into account a wide range of important considerations for and against making a grant, and aggregates these considerations in a thoughtful manner. The most important considerations are listed in LV's "grant investigation" summary, and we think each of these seem highly relevant to an analysis of the grant's marginal cost-effectiveness. We'll discuss the parts of LV's process that seem particularly relevant:

- Focusing on counterfactuality and coordination with other funders
- Having a clear theory of change behind each grant
- Considering the opinions of experts
- Aggregating this information into a decision

Counterfactuality and funging

LV's process involves considering the counterfactual where it doesn't provide funding, and particularly whether another funder would otherwise fund the project. Funding a grant that would otherwise be funded can still be the best thing to do — for example, if it is displacing the funding of another high-impact grantmaker, and there are reasons for LV to be better suited to provide the grant (for example, in being able to provide more support and followup to the grantee). Given the fungibility of some ECF grants, we think what's important here is that LV is successfully coordinating with other impact-focused funders.

While we had limited access to direct examples of such coordination, we have seen a large amount of indirect evidence (for example, reference to situations where LV and another funder agreed to both fund a particular opportunity and worked together to decide the grant amounts, and where LV shared some of its internal reasoning with another funder and vice versa). We are confident that LV is directly in touch with other funders, and is being thoughtful about how to coordinate with them, but we want to caution that we have not directly checked detailed examples of this coordination.

Theory of change

Many of the questions LV fills out in its standard grant investigations template help make it clear what the theory of change is behind the grant. In the cases we saw, we felt like we could follow this theory of change quite clearly, and this made us more confident in the cost-effectiveness of the grant. Importantly, this often came with careful consideration of a variety of important factors, such as:

- The downsides of the grants.
- The track record of the individuals the grant supports.
- Critical assessments of the internal impact evaluations of the grantee organisation.
- Views of various expert advisors on how well the team was executing or how promising their plans were.

Network of advisors

Another point worth noting in LV's process is its extensive network of expert advisors. While we cannot share the details of the list, a much shorter list of advisors is provided on LV's website. We received access to the full list, and perhaps more importantly, saw in many of the grants that these advisors were directly involved in discussing the case for and against particular funding decisions. This gives us confidence that LV is able to provide grants to areas even when it doesn't have the in-house technical expertise to assess the grants itself (though we also noted that LV's grantmaking team seemed generally able to engage with advisors and grantees on many of the technical details).

Aggregating information

Almost all the value of raising key considerations comes from how it informs a particular funding decision — therefore, how LV aggregates its findings is crucial. We were left with a very positive impression of how LV highlighted which of the considerations were most important, why it did so, and how this informed a specific funding decision. Unfortunately, it's difficult to convey examples of this without providing confidential information, but some general features of this include:

- Highlighting the key premises and evidence for each grant.
- Explaining the case *against* the grant, and in the event LV makes the grant, explaining why this case against it was weaker than the case for it.
- Sharing the key outcomes that LV hopes to achieve with the grant (including in many cases making forecasts about these).

LV also uses a comparative process to decide which of its various funding opportunities it should support:

- Creating a rough, informal rank-ordering of the grants.
- Funding the highest-ranked options available (while also considering alternatives, like saving the money and giving later, or sharing the funding opportunity with another funder).

We think this practice seems like a reasonable way to prioritise the most cost-effective funding opportunities (and particularly, is better than an alternative of just deciding each grant on a case-by-case basis, rather than directly comparing the available options).

The cost-effectiveness of Longview's past grants

While we had a positive impression of LV's processes, what matters most is that they lead to cost-effective grants in practice. We'll share two anonymised examples of how the process was applied in practice: one of an ECF grant that we checked in some depth, and one of a grant made by LV that was not part of the ECF.

Starting with the ECF grant:

- The grant was referred to LV by another funder (highlighting the kind of coordination we were interested in). There were seemingly good (but unfortunately, confidential) reasons for the ECF to be best-placed to provide the grant.
- In checking the grant, LV reached out to two independent and qualified advisors (arguably, some of the most qualified people we could have thought of for the particular grant). It also engaged with the previous funder (who also had direct technical expertise) on its understanding of the case in favour of the grant.
- The grant required some assessment of technical details, and we were impressed by:
    - The level of depth with which LV's staff engaged with these details.
    - Their level of engagement with advisors.
    - How clearly this technical information was explained (for example, so that we could follow the reasoning).
- There was attention to grant sizing, including creating a specific plan to follow up on the grant afterwards and providing funding conditional on the performance of the grant in its first year.
- The grant carefully considered the counterfactual of what would occur without funding, and made a clear case that providing funding would be especially valuable.
- The internal reasoning included references on the particular people the grant would fund (which were very positive).

Overall, we could follow the case for providing the grant, and agreed that it seemed plausibly highly cost-effective. There were no significant missing considerations, or problematic inferences that we could find (despite trying). We want to be transparent about this, but also not overstate it — this was a grant in a technical area we do not have direct experience with: all we can conclude is that the reasoning appeared logical and sufficiently comprehensive, but we needed to defer to LV and its advisors for most of the key premises.

We now turn to a similar analysis of a more general grant that LV provided, which was made several years ago. In this case, we felt more confident that we did have the relevant expertise to engage with LV's reasoning (as we were familiar with the grantee), and we were able to check the results of the project that was funded. Some features of this grant included:

- Highlighting and discussing potential conflicts of interest.
- Analysing the grantee's room for funding, taking into account the funding being provided by other funders. This case seemed particularly representative of the kind of fungeing concerns we've highlighted, because multiple impact-focused funders were looking to provide funding.
    - The specific calculation LV used to decide how much funding to recommend was redacted in the documents we saw, but it was explicitly aimed at avoiding adversarial situations whereby multiple funders want the grantee to be funded, but they do not fund it themselves in the hope that another funder will do so.
- Considering the individual track records of the people the grant was funding. One caveat is that though we agreed these were quite positive, no negative points were made, which did suggest some lack of even-handedness.
- Reviewing the grantee's own impact evaluation in a reasonably in-depth way.
- Directly considering the case against the grant, including:
    - Potential downsides of the grantee's activities
    - Reasons the grantee may fail to replicate their past success
- Including concrete forecasts from LV on key metrics for the grantee that it expects to be achieved in the years after the grant, and discussing what it might look like if the grant went poorly. While LV did share with us two followup documents to this grant, it did not comprehensively assess whether its predictions were accurate. However, it was possible to check these forecasts based on publicly available information, and we found that the grantee performed at or above the expectations implied by LV's forecasts.

We assessed several other grants LV had made (both via the ECF and otherwise), and while we found some variation in their comprehensiveness (discussed more below), we overall got a very positive impression of how well LV's processes function, and the expected cost-effectiveness of the resulting grants.

The structure of the ECF

Another reason in favour of the ECF as a recommended fund for donors to support, and as a recipient for our new Risks and Resilience Fund, is that its general scope and structure aligns well with our goals. Our Risks and Resilience Fund is intended to be a default giving option for people with a wide variety of values who are interested in reducing GCRs, and the ECF was set up to be a default giving option for a wide variety of people looking to improve the long-term future. Though improving the long-term future and reducing GCRs are different, most current work on the former is focused on the latter, and we expect that for a significant number of donors, long-term-oriented motivations will be one key reason they are interested in reducing GCRs.

The aspects of the ECF's scope and structure that are most relevant here are:

- **Requiring some degree of understandability to donors.** This has clear upsides, but also does have a downside in that illegible but impactful options may not be considered. We discussed some of the tradeoffs of this response in a comment on the announcement of the ECF (because at the time, we were managing the fund's communications).[6]

- **Supporting a diversity of causes.** We think by supporting a range of causes, the fund is better suited to act as a default option for a wide range of donors who between them have diverse views about which cause might be most promising. The fund discussed its approach to cause prioritisation in its 2022 grant report, including sharing that if one cause ends up being notably and persistently more cost-effective than others, the ECF may no longer require supporting a range of causes and instead just make the most cost-effective grants possible.

The fund values legibility and a degree of pluralism. We think these generally point in favour of the fund, given our goal is having a strong default for donors, but we think it is worth highlighting that these values also act as constraints that could trade off against cost-effectiveness. In particular, the fund is not an ideal fit for those who expect illegible options to be more cost-effective, or have strong views that a specific cause within the cause area of reducing GCRs has the most cost-effective funding opportunities on the margin.

We should also again emphasise our conflict of interest here. We were heavily involved in the launch of the ECF, and we pushed for its creation because we thought it would be valuable to have a fund that could act as a strong default for those looking to give to charity to improve the long-term future. Nevertheless, we think the fact that the fund was created with very similar goals to those we have for our new cause area funds is generally a point in its favour.

Alternatives to the ECF

Though we think the ECF is a promising option for donors looking to reduce global catastrophic risk, one reason we recommend it is that there aren't many alternative options available, and that we haven't looked into all the available options yet.

We discuss how we think the ECF compares to EA Funds' Long-Term Future Fund below. However, as we explain there, we don't yet feel able to directly compare them in terms of their marginal cost-effectiveness for donors, and hence recommend both. In the future, we aim to also look into Founders Pledge's Global Catastrophic Risks Fund, and this may change our view of the ECF and the LTFF.

We don't mean to say there are no other funders in this space. For instance, there is Open Philanthropy, the Survival and Flourishing Fund, Lightspeed Grants, and Manifund. Yet of the options we are aware of, our understanding is that regular donors can only support Manifund, which isn't an expert evaluator but a platform that allows individuals to directly fund or request funding to organisations or regrantors.

[6] One note: Michael Townsend shared in that comment, "I think I may share your intuitions that some of the smaller grants the Long-Term Future Fund makes might be more cost-effective than the typical grant I expect the Emerging Challenges Fund to make (though, it's difficult to evaluate this in advance of the Emerging Challenges Fund making grants!)" in the linked comment. This was based on a much less in-depth understanding of the LTFF and the Emerging Challenges Fund than these evaluator reports are. Now having a richer understanding of both funds, Michael no longer has this view.

We think that there's room for more publicly accessible charity evaluation and grantmaking projects working to reduce GCRs, and would be excited to see such initiatives.

Other relevant findings

In addition to factors directly influencing our decision to recommend the ECF, we want to share some of our other findings relevant to LV and the ECF, which may be relevant to donors, LV, and other evaluators.

Some inconsistency in comprehensiveness of grant investigations

We highlighted two specific grants above that were generally quite comprehensive in the range of considerations that were provided. We found some of the grants LV made were less comprehensive, and there were some inconsistencies, such as:

- Variation in whether LV made forecasts or verified its success (forecasts are an optional part of LV's grant investigations process, and LV shared that it now only verifies its forecasting track record in cases where it renews a grant, but not in general).
- Variation in whether all the considerations highlighted in LV's grants investigations process were made. For example, in at least one grant we saw, there was very limited discussion or consideration for the sizing of the grant.
- Large variation in the extent of the reasoning per grant (at least that was shared with us). We aren't sure whether this was just because LV shared different amounts of its internal reasoning for each of its grants, or whether there actually was significant variation in LV's comprehensiveness behind different grant investigations. We also don't think this is necessarily a negative point, as the extent of reasoning seemed to largely correlate with the size of the grant.

Another point was that LV doesn't have a clear metric it uses to decide how beneficial each grant was, nor a defined "bar" it uses to decide which opportunities to fund (in contrast, several other evaluators we investigated do use such systems to ensure consistency in their approaches, including the LTFF). We thought this was worth highlighting, but after raising this with LV, we were persuaded that this may be the right decision in its context. Our understanding of this is:

- Creating a clear metric, and defining where the "bar" is using that metric (and keeping it up to date as their expectations of the availability of future giving opportunities and reserve of funds change) is a significant amount of work that could instead go into finding more grants or better assessing them.
- Due to its lower volume of grants, it feels it is able to keep track of the relative cost-effectiveness of different funding opportunities considered, and so its approach of rank-ordering them and deciding the bar for funding on a more case-by-case basis is more appropriate.
- LV would consider changing its approach if ts grant volume increased such that this were no longer true.

GWWC's lack of even-handedness in the ECF reports

In almost all cases where we could see the ECF's internal reasoning behind a grant, we saw consideration of possible downsides or the case against making the grant. However, this information was not provided in ECF's public reports. To a large extent, this is attributable to us (GWWC), as we were responsible for managing the fund's communications and writing the grant reports with support from LV. While we think the reports do contain useful information about the grantees, and found LV's quotes (included in grants reports) generally do a good job of concisely describing its internal view, the information the ECF provides is somewhat limited in value for donors looking to clearly understand the thinking that went into the grants. This is especially true compared to the LTFF, which provides less information overall, but does seem more even-handed in the information it does share.

We are not sure what the future of the ECF's communications will be with LV now managing them, but we think it would be worth it for LV to provide a more even-handed treatment of grant reasoning in them than we did in collaboration with LV in the past.

How does the ECF compare to the EA Long-Term Future Fund?

We weren't able to make a useful and justifiable comparison between the ECF and the LTFF in terms of the cost-effectiveness of an extra dollar a donor provides as part of this evaluation. This was due to the general limitations of our approach, and in particular the relative difficulty of comparing across causes and interventions and coming to strong views on the marginal cost-effectiveness of grants in reducing GCRs (as compared to global health and wellbeing and animal welfare).

Given both LV and the LTFF have shown the capacity to grant in the millions of dollars per year and seem to have ample capacity to grant more, we also don't expect donations to our Risks and Resilience Fund to strongly diminish either's capacity to make cost-effective grants in the short term. We therefore currently plan to allocate half of our Risks and Resilience Fund's budget to each, though we could adapt this in the future if our Risks and Resilience Fund donations end up making up a large proportion of their grantmaking.

There are still some differences between the ECF and the LTFF that we think are worth highlighting to donors, so they can come to their own views on which may best align with their values and starting assumptions.

The three most important ones seem to us:

- Allocation of funding by cause:
    - The LTFF's grantmaking contains more AI-risk-related grants than the ECF's grantmaking.
- The types of organisations and grants they consider:
    - The ECF generally funds larger organisations to conduct specific projects, hire more staff, or cover other expenses.
    - The LTFF is more likely to give grants to individuals, or newer and smaller projects.
- Quality of processes:
    - We thought the ECF had less room for improvement on the quality of its processes and its implementation than the LTFF (but we don't think this implies the ECF makes more cost-effective grants).

Some other differences we observed include:

- The level of time that can be spent investigating and writing about each grant:
    - LV makes fewer but larger grants, and is able to spend more time on them, and to explain its reasoning both internally and externally.
    - LTFF generally has a much higher volume of smaller grants, and so puts less time into evaluating and explaining each individually.
- The counterfactuality of grants:
    - In general, giving to the LTFF is more likely to directly fund projects that would otherwise not be funded.
    - The ECF is more likely to displace funding from other funders that try to reduce GCRs, such as Open Philanthropy, or the major donors LV advises. (which, as explained earlier, we don't necessarily see as a negative)

Conclusion

We recommend the ECF to donors wanting to maximise their impact in reducing GCRs, and we plan to allocate half of our Risks and Resilience Fund's budget to it.

In short, this is because:

- LV's grantmaking processes seem suitable for making highly cost-effective grants in reducing GCRs.
- LV's grants seem plausibly highly cost-effective.
- The ECF's structure and scope is relevant and understandable to a wide variety of donors looking to reduce GCRs.
- We don't know of any clearly better alternative donation option in reducing GCRs.[7]

We think ECF has some room for improvement:

- Greater consistency in the comprehensiveness of grant investigations.
- Having more even-handed communications.

By comparison, EA Funds' Long-Term Future Fund had more room for improvement on its grantmaking processes. However, we found no clear, justifiable reasons for the marginal dollar to be better spent at the ECF than at the LTFF (or vice versa). As a result, we recommend both as top-rated funds and plan to allocate half of our Risks and Resilience Fund's budget to each until our next evaluation.

In a future evaluation, we would like to:

- Compare the ECF to more options.
- Investigate more of the ECF's past "hits" (its historically best grants).
- Seek input from external reviewers and experts earlier in the evaluation process.

Doing so would give us more confidence in a decision about whether to continue to recommend the ECF, and to allocate funding to it from our Risks and Resilience Fund.

---

[7] Note that we haven't yet evaluated Founders Pledge's Global Catastrophic Risk Fund, but aim to do so next year.

Acknowledgments

We would like to thank the following people for providing helpful feedback on an earlier draft of this report: Matt Lerner and Tom Barnes (Founders Pledge), and Joshua Monrad (Effective Giving). We are grateful for the time they took to review our work. We would also like to extend our gratitude to Longview Philanthropy and its team for working with us on this project. Last but not least, we would like to thank Katy Moore for copy-editing. All views expressed here are our own, as are any remaining errors or omissions.

Appendix

Our starting questions for 'evaluating the evaluators'

Before investigating any evaluators, we wrote a list of questions that we would consider. We didn't expect to be able to answer every question — rather, we focused on what we expected the biggest cruxes to be — but we wanted to have some kind of template that would help us with our evaluation.

Worldview and approach

1. Who are the beneficiaries they want to maximise positive impact for, and to what extent do they weigh each beneficiary's interests compared to others?
    a. I.e. who do they include in their moral circle and to what extent?
2. What do they consider positive and negative impacts for these beneficiaries?
    a. How do they define and measure positive impact?
3. Are there further worldview-based constraints or principles they take into account in determining their recommendations?
    a. E.g. do they screen for equality/diversity considerations?
4. Do they have any further decision-relevant views/assumptions or do they employ any further heuristics/rules-of-thumb in making their recommendations?
    a. E.g. do they prefer certain types of evidence over others; do they try to account for harder-to-measure and/or longer-term effects; do they have any preferences for "hits-based" or "highly evidenced" recommendations?

Quality of evaluation and grantmaking

1. Are they trying to find the highest-impact funding opportunities (according to some worldview) and acting as such? (**focus on impact**)
    a. Do they claim they are trying to find the highest-impact funding opportunities?
    b. Does their self-stated bar for making a recommendation imply their recommendations are among the highest-impact funding opportunities?
    c. Can/Do they make a plausible case for why their cause area focus is in alignment with finding the highest-impact funding opportunities, given their worldview?
    d. Can/Do they make a plausible case for why their evaluation methodology and approach is in alignment with finding the highest-impact funding opportunities, given their worldview?
    e. Are they continuously looking for improvements so they can get to / keep finding the highest-impact funding opportunities?
        - E.g. do they invite and implement feedback?
        - Do they self-evaluate (e.g. how their past recommendations have panned out) and update as a consequence?
        - Do they learn from mistakes?
2. Do they apply the best-available evaluation methodology and practices? (**rigour/methodology**)
    a. Do they have a sufficiently clear bar for recommendations, and a transparent test for whether funding opportunities meet that bar?

b. How do they sample their causes, interventions and charities?
    - Top-down or bottom-up (e.g. on request)? What heuristics do they use to choose whether to evaluate a cause/intervention/charity?
c. Do they use all types of evidence available?
d. Are they good at integrating evidence to come to an overall judgement / do they apply the right methodology for this?
    - E.g. Bayesian vs frequentist
    - E.g. CEAs, their comprehensiveness, how much they are relied on, how they deal with uncertainty
e. Do they only look at track record and (extrapolating) cost-effectiveness estimates/results, or do they also form an inside view on the theory of change of a charities' programmes?
    - And if so, how well-founded is this, and how does this influence their evaluation?
f. Do they apply strategic grantmaking considerations appropriately?
    - room for funding
    - fungeing and leverage
g. Generally: what may they be (structurally) missing, given the framework/methodology and heuristics they use?
h. Do they engage in external (peer) reviews or seek expert opinions to validate their evaluations?
i. How do they ensure that different perspectives and voices are considered in their assessments?
j. To what extent are there any positive or negative externalities for charities or other actors participating in their evaluation processes?
    - E.g. do they provide useful feedback; do they ask a lot of time; do they communicate clearly; do they generally take a supportive attitude or do they sometimes do things at the expense of a charity; are there risks involved for a charity to participate in their evaluation; do they collaborate and share lessons with other evaluators?

3. Do they have the appropriate in-house resources to implement their stated methodology? (**implementation/resourcing**)
    a. How much staff time goes into (each step off) their evaluations?
    b. Do they have the appropriate in-house expertise?
    c. Do they have access to and utilise the appropriate external expertise?
    d. Do they have the appropriate network, reputation and resources to get access to all the information they need?
        - Do they have access to unique funding opportunities via their network? How do these reach them?
        - Are they able to identify opportunities that others may not, e.g. due to internal expertise?
    e. How often do they update their recommendations? How up-to-date are their recommendations?
    f. How have organisations they've recommended performed over time?
    g. How many mistakes do they make in their evaluations, e.g. in their CEAs?
    h. How many charities and interventions drop out of their process and at which point? What does their “funnel” look like and is it consistent with them selecting the highest-impact funding opportunities?

        - How often do they decide not to recommend a charity they've seriously started investigating?
    i. In how much detail do they check everything?
        - How many things do they check about their recommendations (e.g. do they just check results or also organisational competence, theory of change, etc.)
        - What does their due diligence process look like
        - Do they follow up on claims by orgs (e.g. with third parties), references and citations? Do they double-check things?
        - Do they check how the money granted/recommended is ultimately used by the org they recommend / to what results it leads?
4. Are incentives internally and externally aligned with them making impact-maximising recommendations? (**incentives/conflicts of interest**)
    a. Does the organisation have other interests (such as other teams with their own goals/metrics; fundraising interests) that could get in the way?
    b. How independent is the research team within the organisation?
    c. How are relationships with charities set up?
    d. Are there any downsides/upsides to the organisation to removing/recommending charities?
    e. Are there any other potential conflicts of interest?
    f. How have their recommendations and their processes changed over time, and why?
5. Do they communicate transparently and accurately? (**transparency/scout mindset**)
    a. To what extent do they share their reasoning for recommendations?
    b. Is their public communication generally truthful and nuanced?
        - E.g. do they qualify their statements / avoid making overconfident claims (e.g. "this has x more impact than this")
    c. Do they highlight their uncertainties and the limitations of their approach?
    d. How accessible and user-friendly are their evaluation outputs for different stakeholders, including donors and charities?
    e. Do they acknowledge and learn from their mistakes?
        - How do they react to criticism?
    f. Does their website and other public communications accurately reflect their approach, worldview and quality, based on our evaluation?