

Founders Pledge

Investing to Give

AUTHOR: SJIR HOEIJMAKERS

OCTOBER 2020



Executive summary

In 1784, the French mathematician Charles-Joseph Mathon de la Cour wrote a parody of Benjamin Franklin's then-famous *Poor Richard's Almanack*. In it, Mathon de la Cour joked that Franklin would be in favour of investing money to grow for hundreds of years and then be spent on utopian projects. Franklin, amused, thanked Mathon de la Cour for the suggestion, and left £1,000 each to the cities of Philadelphia and Boston in his will. This money was to be invested



and only to be spent a full 200 years after his death. As time went by, the money grew, and in 1990 Boston received an impressive \$5 million and Philadelphia \$2.3 million, which was spent on charitable causes on behalf of Ben Franklin.¹

Benjamin Franklin is one of the first people we know of who practised *investing to give*: purposely investing funds at one point in time in order to have more impact later. This report investigates how promising this strategy is today, and whether we could do even better than Franklin did. In particular, we are trying to answer whether, if we want to maximise our impact as philanthropists, we should do one of two things: either give to the highest-impact opportunities available now, or invest in order to give even more impactfully at a later date. At Founders Pledge, we are considering launching a Long-Term Investment Fund for our members who would like to invest to give for maximum long-term impact.² This Fund would take contributions from members, invest them, and disburse the resulting funds to nonprofits at those times when the long-term impact of doing so appears highest, whether this is in five years or in 500 years. This research project on investing to give is key to our ongoing decision process on whether we should create such a Fund. Therefore, this project's primary purpose was to evaluate investing to give from a long-term impact perspective, but we have also looked into its potential from the perspective of benefitting the current generation, and from the perspective of averting animal suffering in the near term.

In this summary, we highlight the key findings of our research project and their practical significance. For a more detailed and exhaustive explanation of our approach, our model, and the evidence and reasoning supporting our findings, we refer the reader to our full report.

1. Mary and our proxy model

We start by answering a proxy question, featuring the fictional Founders Pledge member Mary. Mary cares deeply about others regardless of where or when they live. She has \$1 million that she wants to spend on making the world a better place in a way that has the highest expected long-term impact: she is open to opportunities that have a high chance of failing but would yield an outsized reward if successful. So, how can Mary best achieve this impact? Should she allocate her \$1 million to a Fund which invests her money and then gives to the highest-impact funding opportunity Founders Pledge is able to find? Importantly, for this proxy question, we choose to disregard what we call *investment-like giving opportunities*. We also fix the timeline of investing to give at 10 years, and assume equity market index funds as our investment strategy. The implications of releasing these restrictions are discussed later.

In order to answer the question, we estimate Mary's expected impact of investing to give *relative to* her expected impact of giving today. We identify three key factors:

- The financial returns we are able to achieve in 10 years: the financial returns factor
- To what extent we are able to ensure that the funds will be spent on high-impact funding opportunities in 10 years: the **persistence factor**
- 3. The difference in cost-effectiveness between the highest-impact opportunities we are able to find and fund now and in 10 years: the **difference in**

cost-effectiveness factor

In our simple quantitative model, we make estimates for each of these factors and multiply them together to reach a ratio for the expected impact of investing to give 10 years later compared to giving now. Intuitively, this can be seen as starting with an impact of '1' for giving now, and then multiplying this number by all the factors that can grow or shrink it over 10 years of investing:

impact of investing to give in 10 years = impact of giving today * financial
returns factor * persistence factor * difference in cost-effectiveness factor

Our estimates for the factors are mostly based on extrapolations of historical data and expert surveys. They come with many limitations and caveats, which are discussed in detail in the report. The precise model results should hence be taken with a pinch of salt, but the overall takeaways are useful.

1.1. Model results

Projected impact ratio

So, how impactful is Mary's investment likely to be? Figure 1 shows our probabilistic estimates of the impact ratio of investing to give in 10 years compared to giving now. The figure illustrates the uncertainty in our predictions: our 90% credible interval³ runs all the way from 0.1 to close to 40.

Figure 1. Investing to give in 10 years compared to giving today



Impact ratio

Source: Guesstimate model

We estimate that Mary will have more impact by investing than by giving today with 70% probability. More importantly, we estimate that the *expected value of the impact ratio*⁴ is very high: Mary will have nine times as much impact by investing *on average*.

The high expected value illustrates a noteworthy asymmetry: Mary has more to gain by investing to give than she has to lose. In the worst cases, her invested \$1 million will have no or negligible impact after it is spent 10 years from now. In the best cases, however, her invested \$1 million could end up having an impact *many times* larger than it would now. These potentially very large gains are much more significant than the potential losses, and drive up the expected impact of investing to give.

1.2. Financial returns

Can we make significant returns on the money we invest? We find that, through equity market index investing, it has historically been possible to quite reliably make positive financial returns on investments, as illustrated by the S&P index values that are charted in figure 2. There is certainly variance, but in the typical scenario one could double one's financial resources – in nominal terms – over the course of a decade. This is why financial returns are the key factor driving our estimate that in the majority of cases (70%), investing to give in 10 years will have a higher impact than giving now.



Figure 2. S&P 500 Index in the past 50 years (nominal)

Source: macrotrends.net

1.3. Persistence

We find that our second factor, persistence of funds and values, has the least overall influence on our estimates. There are meaningful risks of value drift and loss of ownership, but these can mostly be mitigated by a well-designed Fund. Such a Fund could, for instance, avoid risks of value drift by being legally compelled to disburse funds to the charitable sector, and avoid risks of loss of ownership by being legally decoupled from Founders Pledge.

1.4. Difference in cost-effectiveness

When deciding where to give, you want to ensure that for every dollar you donate, you're creating the maximum possible impact: you want your donation to be *cost-effective*. We find the difference in cost-effectiveness of giving later to be the most important argument for investing to give overall, but also the most uncertain one. This is largely due to *exogenous learning*: learning that occurs over time regardless of Mary's giving, and which will allow Founders Pledge to identify higher-impact funding opportunities. Because the field of long-term impact research is so new, there is likely still a lot to learn, even in a matter of years.

According to most of the experts we surveyed, exogenous learning over the next 10 years is the strongest factor driving up the expected cost-effectiveness of the funding opportunities that we would give to in 10 years. They predict it will outweigh other factors that could drive down cost-effectiveness - such as a larger community of impact-focused funders - with high probability. And many think there is a meaningful probability that these learnings could lead to a more-than-tenfold increase in cost-effectiveness. Hence, whilst not driving the outcome in the majority of cases, the difference in cost-effectiveness factor is the main driver for our high *expected* impact ratio.

However, our surveyed experts were not unanimous in this view, there were important limitations to our survey methodology and sample, and we think there is also a somewhat plausible case for a *decrease* in the cost-effectiveness of the best available funding opportunities over time - in particular due to more impact-focused funding outweighing the effects of exogenous learning.

1.5. Near-term aims

We also evaluated our proxy question from the perspective of benefitting the current generation and averting animal suffering. For this, we used input from impact-focused funding opportunity evaluators <u>GiveWell</u>, <u>Farmed Animal Funders</u> and <u>Animal Charity</u> <u>Evaluators</u>.

For benefitting the current generation, we estimate an expected 10-year impact ratio for investing to give compared to giving today of 2.1, and for investing to be higher-impact than giving now in 50% of cases. The difference with the long-term impact case is probably best explained by a lower projected exogenous learning rate: research into funding opportunities in this space is easier to do well, has already received a lot more time and resources, and is supported by a wealth of intervention research from development economics, which has been around even longer.

For averting animal suffering, we estimate an average impact ratio of 4.2, and for it to be higher-impact to invest to give in 10 years in 60% of cases. However, we expect there to be a large amount of investment-like giving opportunities available in this space, which strengthens the case for giving now.

2. Further considerations

2.1. Investment-like giving opportunities

Our proxy question intentionally excluded investment-like giving opportunities: giving opportunities whose primary route to impact is making more financial or human resources available to be "spent" on the highest-impact opportunities at a later point in time. A straightforward example is encouraging people with similar values to invest to give rather than not give at all.

We think these types of opportunities could in principle have similar or higher compound returns than investment, even in the longer term, mainly because they can benefit from exogenous learning as much as investing to give does.

2.2. Option value

The other significant limitation of our proxy question – and of Ben Franklin's approach – is that it assumes disbursements of funds after a fixed amount of years. Usually, when choosing to invest to give, we don't need to commit to any particular timeline of disbursement: the only option we lose is to give now. However, if we choose to give now, we lose all other options. Investing hence has more option value than giving now does. This has three implications.

It is, first of all, relevant for interpreting the 10-year impact ratio estimate in our proxy model: if the estimated ratio is higher than 1 - as it is - this suggests that investing is the highest-impact strategy available at this point in time, excluding investment-like giving opportunities. However, if the ratio were lower than 1, this would not immediately imply that investing is not the highest-impact strategy, as there could be other years for which the ratio would be higher than 1.

Secondly, investing to give allows us to use the "watch then pounce" strategy.⁵ Even if our ability to have an impact were to decrease over time *on average*, it might be that, every once in a while, unusually high-impact funding opportunities arise. When investing, we have the option to fund such opportunities, whereas giving now restricts us to the funding opportunities available right now.

Thirdly, investing to give allows us to learn more about the question of when to give itself, before fixing the timing of disbursement of our funds.

2.3. Higher-return investment strategies

For our proxy question, we have assumed equity market index funds as our investment strategy. There are alternatives that might yield higher expected returns, such as leveraged index investing or venture capital investing. However, due to restrictions on higher-investment-risk investments by nonprofits in many countries, it might be difficult to set up a charitable investment fund that pursues these strategies for a large part of its investments.

We will consider both these alternatives for future research and in our plans for future investment strategies of a potential Long-Term Investment Fund.

2.4. Others' time preference

Do most people care more about the present or the future? There is a wholly separate, theoretical argument for investing to give, rooted in the hypothesis that many people have an intrinsic preference to receive benefits now over receiving them in the future. To what extent this is true is a matter of substantial academic debate,⁶ but most economists seem to agree that there is very likely some positive rate at which most people intrinsically discount future benefits.⁷

Other people's focus on the present makes it more plausible that investing to give is a high-impact strategy in two distinct ways. First, it helps explain why we can obtain large returns on financial investments: if most people value the future less than us, they will be willing to trade their influence over the future for more influence over the present. And they will leave to us some of the investment opportunities that are advantageous from the perspective of benefitting the long term. Second, it makes it more likely that humanity as a whole is overspending in the present, and that by offsetting this, we can improve humanity's overall spending portfolio by investing.

2.5. Considerations found to be less important

In the report, we address many other considerations which have some weight in our overall view of investing to give. These include general considerations on compound returns on giving now, investment and fund management costs, the impact of investing itself, and correlations among our key factors. We found none of these to be strong enough in either direction to affect the overall balance, though they do increase our overall uncertainty.

3. Conclusions and recommendations

3.1. Benefitting the long term

Based on the analysis in this report, we think investing to give is a very promising strategy for any long-term-oriented individual philanthropist to consider at this point in time. This is primarily because of current opportunities for exogenous learning and financial returns, and because of the option value investing to give carries. However, giving to investment-like giving opportunities could be a good alternative to investing to give.

Given the many limitations in our analysis, we do not think we can confidently claim investing to give is a higher-impact strategy than giving today. Still, we believe it is plausible that it could be, and we recommend our members to at least consider making it a part of their philanthropic portfolio.

For Founders Pledge in particular, the conclusions of this research project are a strong argument to set up a Long-Term Investment Fund for our members. We have started exploring the practicalities of doing this, and tentatively expect to launch such a Fund in 2021.

3.2. Benefitting the near term

We think investing to give is an option worth considering for members focused on benefitting the current generation, but only when a suitably designed investment and granting vehicle becomes available. We think it is less promising than for members focused on benefitting the long term, mainly because there seem to be fewer opportunities for exogenous learning.

Similarly, we think investing to give is an option worth considering for members focused on averting animal suffering in the near term. However, we wouldn't be surprised if there are investment-like giving opportunities in this space that can outperform or at least match investing to give. We are generally less confident in our conclusions here than in the case of benefitting the long term or the current generation of people.

3.3. When should we stop investing and start giving?

If investing to give is indeed a higher-impact strategy than giving now - at least from a long-term impact perspective - then when is this no longer the case? When should we stop investing and start giving? Our model and analysis show multiple ways in which this could occur. For instance, we could see the exogenous learning rate diminish - and not expect it to go up again - as to some extent seems to have already happened in the case of benefitting the current generation. Or expected investment returns could go down substantially. Also, there could be times with extraordinary funding opportunities during which we should give at least part of our funds (the "watch then pounce" strategy).

Lastly, our recommendations here target the individual philanthropist and become less valid at very large (>\$100 million) philanthropic budgets. At that point, whether you immediately give or invest to give starts to meaningfully influence the distribution of high-impact giving and investing at a global level, and new considerations come into play, such as diminishing marginal returns to giving in a particular year. Furthermore, the "watch then pounce" strategy arguably has diminishing marginal returns: most of its value can be achieved by a Long-Term Investment Fund with a limited amount of funds.

Acknowledgements

For various contributions to this report, including discussion, advice and expertise, we are grateful to:

- Philip Trammell and William MacAskill, Global Priorities Institute, Oxford University
- Carl Shulman and Max Daniel, Future of Humanity Institute, Oxford University
- Robin Hanson, George Mason University
- Benjamin Todd and Howie Lempel, 80,000 Hours
- Kit Harris, Longview
- Patrick Kaczmarek, Effective Giving
- Holden Karnofsky, Open Philanthropy
- Alex Cohen, GiveWell
- Kieran Greig and Mikaela Saccoccio, Farmed Animal Funders
- Jamie Spurgeon, Animal Charity Evaluators
- Michael Aird, Convergence Analysis
- Erik Byrenius, Trellis Road
- David Langer, Lionheart Ventures
- Michael Dickens, independent contributor
- Catherine Roberts, independent contributor
- Phil O'Riordan, independent contributor
- Sasha Cooper, independent contributor
- Grayden Reece-Smith, independent contributor

We are especially grateful to Philip Trammell, Carl Shulman, Robin Hanson, Benjamin Todd,

Kit Harris, Michael Aird, David Langer, Michael Dickens, Phil O'Riordan, Sasha Cooper and

Grayden Reece-Smith for reviewing a draft of this report and/or earlier versions. Any mistakes are our own.

Contents

| Introduction | 20 |
|---|----|
| 1. Our approach | 23 |
| 1.1. Why not diversify? | 23 |
| 1.2. Taking the individual philanthropists' perspective | 24 |
| 1.3. Starting from a proxy question | 25 |
| 1.4. Focusing on decision-relevance | 27 |
| 2. The impact of investing to give in ten years | 28 |
| 2.1. The model | 29 |
| 2.2. Financial returns | 30 |
| 2.3. Persistence of funds and values | 31 |
| 2.4. Difference in cost-effectiveness | 33 |
| 2.5. Impact ratio | 36 |
| 2.6. Benefitting the current generation | 41 |
| 2.7. Averting animal suffering | 46 |
| 3. Further considerations on investing to give | 51 |
| 3.1. Considerations favouring giving now | 52 |
| 3.2. Considerations favouring investing to give | 57 |
| 3.3. Considerations increasing uncertainty | 61 |
| 4. Conclusions and recommendations | 65 |
| 4.1. Benefitting the long term | 65 |
| 4.2. Benefitting the near term | 67 |
| 4.3. Suggestions for future research | 68 |
| Appendix I: Financial returns estimate | 72 |
| Appendix II: Persistence estimate | 76 |
| Appendix III: Survey questions | 82 |

| Appendix IV: Analysis of survey results | 93 |
|---|-----|
| Endnotes | 110 |

Introduction

In 1784, the French mathematician Charles-Joseph Mathon de la Cour wrote a friendly parody on Benjamin Franklin's then-famous *Poor Richard's Almanack*, titled *Testament de M. Fortuné Ricard (Testament of Fortunate Richard*).⁸ The main character in the parody, Fortunate Richard, leaves five lots of hundred *livres* (the French currency at the time) in his will with the aim of turning this into a fortune for good through compound interest. After 100, 200, 300, 400 and 500



years each, the resulting billions and trillions of livres are to be spent on utopian projects.

The story reached Ben Franklin, who wrote Mathon de la Cour to thank him for his inspirational idea.⁹ And so Franklin decided to bequeath £1000 (worth roughly \$160,000 today)¹⁰ each to the cities of Boston and Philadelphia, which was to accrue interest over 100 years, then to be partially spent for the cities' benefit, then to be reinvested for another 100 years, and only then to be spent fully. Franklin himself projected that his legacy would amount to more than £2 million for each city 200 years after his death.¹¹

Although his plan probably didn't work out exactly as he would have hoped, it did roughly succeed: 100 years after Franklin's death in 1790, Boston's funds had grown to \$391,000, of which \$100,000 was reinvested and the rest was largely spent on establishing what is now the Benjamin Franklin Institute of Technology. Philadelphia's investments had been less successful: it had stuck closely to Franklin's will's original investment strategy of loans rather than investing in stocks, and as a consequence had only grown to \$172,000. \$39,000 of this was reinvested, and the majority of the rest was given to a science education museum. After 200 years, in 1990, the reinvested sums had grown to \$5 million and \$2.3 million respectively. Boston again gave its money to the Benjamin Franklin Institute of Technology, and Philadelphia spent the money on a variety of causes throughout Pennsylvania, including scholarships.

Benjamin Franklin is one of the first people we know of who practised 'investing to give': purposely investing funds at one point in time in order to have more impact later. He did so quite successfully,¹² and his story carries lessons for how to do even better. In particular, in Franklin's case, three things were fixed that ultimately constrained his impact:¹³

- The beneficiaries of his giving: the funds had to be spent on Boston and Philadelphia, which probably weren't the places where the money could have been used most impactfully at the time of disbursement.
- The investment strategy: the types of investments to be made were to some extent fixed, which meant Boston and particularly Philadelphia weren't able to grow the funds as much as they could have otherwise.
- The timing of disbursement: the funds had to be disbursed at 100 and 200 years after Franklin's death, which were unlikely to happen to be the times at which they could have the largest impact.

What if we were to invest to give today, and approach it strategically? The purpose of this report is to investigate how promising that is. In particular, we are trying to answer whether, if we want to maximise our impact as philanthropists, we should give to the

highest-impact opportunities available now, or invest in order to give more and to higher-impact funding opportunities at a later point in time when it seems most impactful.

At Founders Pledge, we are considering launching a Long-Term Investment Fund for our members who would like to invest to give for maximum long-term impact.¹⁴ This Fund would take contributions from members, invest them, and disburse the resulting funds to non-profits at those times when the long-term impact of doing so appears highest. This research report is an important input in our ongoing decision process for such a Fund.

1. Our approach

1.1. Why not diversify?

A natural question that might arise is: why either/or? Why not invest to give part of our resources and give away the rest now?

A philanthropist may have personal reasons for strategy diversification. For instance, before deciding on where to give most of your funds, you might want to spend some time learning about different worldviews and causes. And while doing this, you might prefer to already spend some of your funds in the causes you are exploring.

From an impact perspective, however, it is likely that either fully investing to give or giving now is optimal for you in any particular year. This is because, where giving is concerned, you are a small part of a global community of philanthropists with similar goals.¹⁵ It makes sense for this community as a whole to diversify its strategy, i.e. to spend a percentage of resources now and invest the rest, and there will be an optimal portfolio choice at this community level. The objective for each individual should be to either fully give now or invest to give so that the distribution at the community level moves closer to this optimal portfolio.

The only cases in which an individual or organisation should diversify to have more impact are (1) when they have unique knowledge of and access to extraordinary funding opportunities, (2) when they control a significant part of the resources that are spent or invested to be spent on certain aims,¹⁶ and (3) when the distribution in the rest of the community is already (very close to) optimal. (1) is certainly possible, but in most cases there is wide access to funding opportunities and there are reasons to share knowledge with other people and organisations with similar aims.¹⁷ For the purposes of this report, we will not be concerned with (2), as we do not expect a single Long-Term Investment Fund to reach such a level in the very near future.¹⁸ (3) seems unlikely to be the case, as there do not seem to be any coordination mechanisms or market forces that would push the community portfolio (fully) towards the optimum. Instead, there are reasons to expect the community portfolio to be biased in either direction for reasons unrelated to impact, and there is no reason to expect these biases will balance out exactly. For instance, people might be overspending because they like to see their funds being used in their lifetime, or they might be underspending because they don't have enough capacity to detect high-impact funding opportunities to spend the funds on.

1.2. Taking the individual philanthropists' perspective

This suggests that one way to approach our question would be to try to calculate the optimal community portfolio for each set of broadly similar philanthropic aims, and then see in which direction current giving differs from it.¹⁹ In this report, however, we choose to take a different approach: we directly compare the expected impact of giving now with the expected impact of investing to give, from the perspective of an individual philanthropist. This is for three reasons:

 Estimating the optimal and current community portfolio is very difficult to do in a way that yields decision-relevant information: it is tricky to define the boundaries of the community, and estimating the optimal community portfolio requires us to make estimates of and/or assumptions about extremely uncertain factors such as the availability of opportunities in the very long term, and the structure of our collective diminishing marginal returns function to spending in any particular year.

- As explained above, we think it unlikely that the choices of individual philanthropists or a new Fund will change whether the community as a whole is overspending or underspending, at least in the near term: this makes considering the community perspective less important.
- Even if we were able to figure out whether a particular community is currently
 overspending or underspending, this wouldn't fully answer our practical question:
 we also need to know whether the community is overspending or underspending to
 a larger or lesser extent than it will be in the future, which would require further
 estimates and/or assumptions on how the community's spending will evolve.

1.3. Starting from a proxy question

Rather than directly addressing the expected impact of investing to give compared to giving now, we first ask a proxy question that is more feasible to answer: what is the expected impact of investing and giving *in ten years from now* compared to the expected impact of giving now?²⁰

In addition, when answering this question, we temporarily disregard what we will call "investment-like" giving opportunities: giving opportunities whose primary route to impact is making more financial or human resources available to be "spent" on the highest-impact opportunities at a later point in time.²¹ We do this because funding investment-like giving opportunities is in some ways more similar to investing to give than to funding non-investment-like giving opportunities. For instance, both investing to give and investment-like giving opportunities allow one to benefit from future learnings on identifying the highest-impact opportunities.²²

Lastly, we will assume equity index investing as our investment strategy.²³ We make this assumption because, relative to most other types of investment returns, there is more evidence available and less controversy about the historical returns on equity index funds and their variance.

On the one hand, answering the proxy question will make us underestimate the potential of investing to give: the case for it is stronger when the timing of the disbursement is not constrained to exactly ten years from now, and a Founders Pledge Long-Term Investment Fund certainly would not restrict itself to such a timeline. Moreover, we might be able to obtain higher returns with other types of investment strategies. On the other hand, the answer to our proxy question does not immediately generalize to a comparison between investing to give and all current giving opportunities, because we are excluding investment-like giving opportunities.

In Chapter 2, we attempt to answer the proxy question. We do this mainly from a long-term impact perspective,²⁴ but we also briefly address it from the perspective of benefitting the current generation and averting animal suffering in the near term.

In Chapter 3, we discuss the further relevant considerations in favour and against investing to give, including the implications of releasing the constraint that the funds will be spent in ten years exactly, compound returns on giving and investment-like giving opportunities, and the availability of other investment strategies. Finally, in Chapter 4, we combine all considerations to come to an overall conclusion on the potential of investing to give as a high-impact philanthropic strategy.

1.4. Focusing on decision-relevance

Throughout this report, we focus on ballpark and practically relevant estimation rather than precision: our goal is to come to practical insights rather than find exact answers to questions. We use many shortcuts, heuristics and approximations. Whether these are technically correct by some definition matters less to us than whether they help us make better decisions. Similarly, the numerical estimates in this report are not meant to express all-things-considered views themselves but are meant as rough but useful approximations to inform our conclusions and recommendations.

2. The impact of investing to give in ten years

We will start by attempting to answer our proxy question from a long-term impact perspective. For this, we consider the fictional Founders Pledge member Mary. She cares deeply about others regardless of where or when they live. She has \$1 million that she wants to spend on making the world a better place in a way that has the highest expected long-term impact, so she is open to opportunities that have a high chance of resulting in little but yield an outsized reward if successful.

We will estimate what Mary's expected impact would be if she allocates the \$1 million to a Founders-Pledge-managed Fund which invests her funds for ten years in equity market index funds and then gives to the highest-impact funding opportunity Founders Pledge is able to find, *relative to* her expected impact if she gives to the highest-impact funding opportunity she is able to find now, excluding investment-like giving opportunities.²⁵ To do this, we look into three key factors:²⁶

- 1. The financial returns we are able to achieve in ten years
- To what extent we are able to ensure that the funds will be spent on high-impact funding opportunities in ten years
- 3. The difference in cost-effectiveness between the highest-impact opportunities we are able to find and fund now and in ten years

2.1. The model

In our simple quantitative <u>model</u>, we make estimates for each of these factors and multiply them together to arrive at a ratio for the expected impact of investing to give ten years later compared to giving now.²⁷ Intuitively, this can be seen as starting with an impact of '1' for giving now, and then multiplying this number by all the factors that can grow or shrink it over ten years of investing:

impact of investing to give in ten years = impact of giving today * financial returns factor * persistence factor * difference in cost-effectiveness factor

For example, say we think our funds will double in size in those ten years, and there is negligible risk of them not being spent on high-impact funding opportunities. However, we think the highest-impact funding opportunity that we will be able to find and fund in ten years will be only three-quarters as cost-effective as the highest-impact opportunity available now. In this case, our financial returns factor will be '2', our persistence factor will be '1', and our difference in cost-effectiveness factor will be '0.75', leading to an expected impact ratio of 2 * 1 * 0.75 = 1.5: we would expect investing to give in ten years from now to have one-and-a-half times as much impact as giving now.

To arrive at a measure of uncertainty, we estimate 90% credible intervals²⁸ for each of the factors. We model the financial returns and difference in cost-effectiveness factors using a log-normal distribution, which is standard for modelling phenomena that can be conceptualized as the multiplication of many independent effects over time.²⁹ This distribution is suitable because it only takes positive values and is heavy-tailed,³⁰ which allows for the possibility of large financial returns or improvements in cost-effectiveness

over time. We model the persistence factor using a beta distribution³¹ because it allows for flexible modelling on the [0,1] interval, to which the persistence factor is restricted by definition.

We use the probabilistic modelling software <u>Guesstimate</u>, which applies what is called Monte Carlo Simulation³² to combine our probability distributions: it draws a random sample of 5000 numbers from a probability distribution to approximate it, and then uses those numbers for subsequent calculations. This means it has some approximation error, but we are able to re-run the simulation multiple times to see whether that meaningfully influences our estimates, and if so, correct for it.

Below, we briefly explain each of the key factors in the model in turn, and give our estimates. The analyses supporting each of those estimates can be found in the appendices.

2.2. Financial returns

The most straightforward advantage of for-profit investing is that it could exponentially grow Mary's financial resources, allowing her to do a lot more good at a later time point (everything else equal).

We will assume Mary's \$1 million will be directly invested into equity market index funds. As explained in the Introduction, we made this assumption because, relative to most other types of investment returns, there is more evidence available and less controversy about the historical returns on equity index funds and their variance. However, we discuss the implications of using other, potentially higher-return investment strategies in Sections 3.2.2. and 3.2.3.

2.2.1. Estimate

Based on historical return rates, we roughly estimate the ten-year financial returns factor to be 2.3 in expectation, with a 90% credible interval of [0.73, 5.0] and a median of 1.9. Equivalently, we expect Mary's \$1 million to grow to \$2.3 million in expectation, to be between \$700,000 and \$5 million in 2030 in 90% of cases, and to grow larger than \$1.9 million half of the time. These figures are all expressed in nominal terms. Please see Appendix I for an explanation of how we arrived at these estimates and why we chose to express them in nominal rather than real terms.

2.3. Persistence of funds and values

Mary's expected financial gains will not lead to more impact if her money is lost before it can be spent (for other reasons than investment losses) or if it is spent on less valuable activities. This could happen through (1) value drift or (2) loss of ownership.

2.3.1. Value drift

Value drift refers to the change of values of a person or organisation over time.³³ For the purpose of this report, we take it to mean a negative change, i.e. one that reduces Mary's positive impact by not giving as effectively or not giving at all. We have included the positive case of value change – value learning – in our definition of exogenous learning under the differences in cost-effectiveness factor.

As the highest-impact funding opportunities may be many times as cost-effective as the average funding opportunity,³⁴ we treat the case in which the funds are disbursed to funding opportunities that aren't chosen based on their high impact as equivalent to the case in which they are not disbursed at all.

Because Mary is using a Founders Pledge Investment Fund to invest and eventually disburse her money, we will examine value drift as it pertains to the Fund rather than to Mary herself. The legal set-up and governance of such a Fund would be chosen with the aim of minimizing risks of value drift.

2.3.2. Loss of ownership

Loss of ownership covers all cases in which we aren't able to allocate Mary's money to the highest-impact funding opportunities available because it is no longer under our control. This could happen, for instance, through a global catastrophe, a change in government policy, or expropriation via legal challenges.

Again, as Mary uses a Founders Pledge Investment Fund, we will examine the risk of loss of ownership as it pertains to that Fund.

2.3.3. Estimate

Based on a short analysis of the various risks, we estimate a 90% credible interval for the persistence factor of [0.8, 0.98] and a median of 0.92. Fitting this to a beta distribution yields an average persistence factor of 0.91. Please see Appendix II for an explanation of how we arrived at these estimates.

We can interpret this as the expectation that a Founders Pledge Long-Term Investment Fund will be able to allocate Mary's invested resources to the highest-impact opportunities available and detectable in ten years from now slightly more than 90% of the time. The main risks driving this estimate down from 1 were those of value drift and loss of ownership due to catastrophe: we expect both to be in the order of a few percentage points over the coming ten years. We have lowered and widened our estimate for the persistence factor further to account for unforeseen risks.

2.4. Difference in cost-effectiveness

The highest long-term-impact funding opportunity we are able to find and fund today will likely differ from the highest long-term-impact funding opportunity we are able to find and fund ten years from now, and this could meaningfully affect Mary's impact from investing to give.

Conceptually, it is useful to separate the difference in cost-effectiveness factor into

- 1. The changing availability of high-impact funding opportunities over time
- 2. Our changing ability to detect those opportunities

2.4.1. Availability of opportunities

The value of the best available funding opportunities changes over time. This value could either increase, for instance via the arrival of new high-impact initiatives that require funding, or decrease, for instance by more funding gaps being filled due to the presence of more impact-focused philanthropists.

2.4.2. Exogenous learning

An advantage of investing to give is that it will allow us to learn about better giving opportunities over time, for instance through improvements in research methodology to detect such giving opportunities.

We should distinguish between two forms of learning – endogenous and exogenous:

Endogenous learning is learning that Mary brings about herself with her giving, for instance by funding research that helps prioritise amongst causes or by funding experiments with new interventions within a cause. Opportunities for endogenous learning can be a reason for her to give now rather than to invest to give.³⁵

Exogenous learning is learning that occurs regardless of Mary's giving. It includes advances in the scientific community, new philanthropic interventions being invented and/or tried out by others with similar aims, moral progress,³⁶ and more. It also captures the time needed for relevant knowledge to become available, e.g. an experiment might take time, research might need to be done in a certain order, or there might be a talent constraint in a research area that takes time to be resolved. When learning is done exogenously, there are advantages to waiting and hence to investing to give.

When Mary puts her funds in the Founders Pledge Investment Fund, she is able to benefit from the long-term-oriented research Founders Pledge does over the next ten years, and from the learnings the Founders Pledge research team derives from external research in that time period. Both of these constitute exogenous learnings from Mary's perspective.

2.4.3. Estimate

The difference in cost-effectiveness factor is by far the most difficult of the three factors to estimate. Instead of doing a full review of all the reasons in either direction³⁷ and weighing all those reasons up ourselves, we decided to survey 11 experts and semi-experts for their all-things-considered views. This gives us a better overview of the diversity of opinion on this question and the uncertainty we should hence have, and it avoids the pitfall of relying too heavily on one particular type of analysis or perspective.

Please see Appendix III for an outline of the survey questions, and Appendix IV for our full analysis of the survey results.

The most important takeaway from the survey is that most respondents think that marginal cost-effectiveness - from the perspective of benefitting the long term - is more likely to increase or stay constant than to decrease over the next ten years and that large increases are more likely than large decreases.

According to most respondents an increase will be largely driven by exogenous learning, but they also mention a few reasons to expect the availability of opportunities to increase, such as new human resources entering the long-term impact space. The main counteracting force mentioned is a decrease in the availability of opportunities due to more funding becoming available over time combined with diminishing marginal returns to spending at the community level. Most respondents predict that exogenous learning and other factors will outpace this decrease, but one expert thinks it will be the other way around. The differences between the respondents' estimates appear to be largely explained by how they weigh these considerations against each other. Based on the survey results, we estimate a credible interval of [0.09, 17] with a median of 1.2 and an average of 4.5 for the difference in cost-effectiveness factor.³⁸ This means that even though in 50% of the cases we predict the cost-effectiveness of the marginal funding opportunity in 2030 to be less than 1.2 times as large as that of 2020, on average we expect it to be 4.5 times as large.

2.5. Impact ratio

The table below shows the bounds of our 90% credible intervals, median estimates and average estimates for each of the factors.

| Ten-year factor | 5th percentile | median | 95th percentile | average |
|---|----------------|--------|-----------------|---------|
| Financial returns | 0.73 | 1.9 | 5.0 | 2.3 |
| Persistence | 0.80 | 0.92 | 0.98 | 0.91 |
| Difference in cost-effectivenes s | 0.09 | 1.2 | 17 | 4.5 |

Combining these factors and their uncertainties, our simple model finds the following

ballpark estimates for the impact ratio:

| Ten-year factor | 5th percentile | 30th percentile | median | 95th percentile | average |
|-----------------|-------------------|--------------------|--------|--------------------|---------|
| Impact ratio | 0.1 | 1 | 2 | 40 | 9 |

The below output from a single run³⁹ of the model in Guesstimate represents the distribution visually:
Projected impact ratio

Showing the 5th - 95th percentile of the simulated probability density function



Source: Guesstimate model

Taken at face value, these results suggest that Mary will on average have 9 times as much impact by investing her \$1 million and giving in ten years than she will have if she gave the funds to the highest-impact opportunities she can find now. They also suggest a high level of uncertainty and/or variability: half the time, Mary is expected to have less than two times as much impact when investing to give in ten years as when giving now, there is a 30% probability that Mary will have *less* impact by investing to give in ten years than by giving now, and there is a 5% probability that investing to give in ten years will yield less than one-tenth of the impact of giving now. On the other hand, we estimate a 5% probability that Mary will have more than forty times the impact if she invests to give in ten years rather than gives now.

Even though the average or *expected* impact ratio is our most decision-relevant estimate in principle, we would caution against putting too much weight on the specific estimate of '9' there. This estimate is highly sensitive to our specific average estimate for the difference in cost-effectiveness factor, about which we are very uncertain. The takeaway that this average estimate is larger than 1 is much more robust, especially as our median estimate is also larger than 1.

Still, the large difference between the average and median illustrates a noteworthy asymmetry: Mary has more to gain by investing to give than she has to lose. In the worst cases,⁴⁰ her invested \$1 million will have no or negligible impact after it is disbursed ten years from now, whether due to losing the funds, value drift, or the unavailability of high-impact opportunities. In the best cases, however, her invested \$1 million can end up having an impact *many times* as large as it could have had when given now, because of large financial returns, because much-higher-impact opportunities become available, because we learn about new crucial considerations in identifying high-impact opportunities, or, most importantly, because of a combination of these: a combination of multiple unfortunate developments will not reduce her impact below zero, whereas a combination of multiple fortunate events can drive it up to tens of times the impact she is able to have by giving now. Potential gains will hence weigh much more heavily on the mean outcome than potential losses. All-in-all, our model suggests investing to give in ten years would be a somewhat risky but potentially very high-impact venture for Mary. As Mary is looking for a way to have the highest expected impact, it suggests she should invest to give, save the further considerations we will discuss in Chapter 3.

2.5.1. Sensitivity analysis

To some extent, our model already incorporates a sensitivity analysis, because we have included our subjective uncertainty around each of the parameters in our estimates. However, it is worth questioning the assumptions we have made and data we have used to arrive at those estimates further to evaluate how robust our conclusions are.

For instance, we have only considered historical data to make our financial returns estimate, and there could be reasons to think the next decade will see lower returns than average.⁴¹ Similarly, our survey for the difference in cost-effectiveness factor has limitations, such as a risk of selection and nonresponse bias, which we have tried to address in making our estimate, but which might still bias our estimate in either direction.

The relevant question then is: what would one have to believe in order for our most important conclusion to change, i.e. that the *expected* impact ratio is larger than 1?

Looking first at each individual factor estimate, keeping the other two constant, it is clear that our conclusion is mostly insensitive to our estimates for the financial returns and persistence factor: even assuming the 5th percentile estimate for each of these factors does not change it. In fact, this is true even for these factors combined: assuming *both* to be at their 5th percentile levels (which our model - which assumes independence - would give only a 0.0025 probability) the estimated expected impact ratio would still be approximately 0.73 * 0.80 * 4.5 = 2.6.

The conclusion is, however, sensitive to our difference in cost-effectiveness factor estimate. Keeping the other two factors constant, assuming our 5th percentile estimate there would lead to an expected impact ratio of approximately 2.3 * 0.91 * 0.09 = 0.19. It is worth noting that this 5th percentile estimate corresponds to the median estimate for one of the experts we consulted, suggesting it may be a defensible position to hold, though it differed strongly from that of all other experts and other survey respondents.

The corresponding reasoning is that - in most cases - more funding becoming available and strongly diminishing returns at the community level will heavily outweigh any exogenous learning or increase in the availability of opportunities in the next decade. In addition, this view holds that large increases in marginal cost-effectiveness due to exogenous learning are very unlikely, or that those only occur in cases in which current marginal cost-effectiveness is extremely low: then they could be compensated by other cases in which current marginal cost-effectiveness is much higher and declines over time.

This does seem like a view one can defensibly hold. However, given the many other defensible views that suggest increasing marginal cost-effectiveness over time, it does not change the direction of our overall conclusion for Mary, again save the further considerations that we will discuss in Chapter 3.

2.6. Benefitting the current generation

In this section, we have a short look at how the potential of investing to give in ten years might differ if your aims are not to benefit the long-term future, but if you are only concerned with benefitting the current generation.

Among the three key factors in our model, the only one that will meaningfully differ is the difference in cost-effectiveness factor: the financial returns and persistence of funds will be largely⁴² independent of what you aim to spend those funds on. We will hence try to get a sense of the difference in cost-effectiveness factor for both aims.

For benefitting the current generation, we have much more direct evidence on the cost-effectiveness of various funding opportunities and how they have changed over time. This is in large part thanks to our research partner GiveWell,⁴³ who have been doing high-quality research into funding opportunities in the global health and development space since 2007 and have been transparently reporting on their findings.

We will first discuss a case study - based on GiveWell's research - to get a sense of how investing to give in 2010 in order to give today could have compared to giving in 2010. Then, we will discuss GiveWell's current views on the future difference in cost-effectiveness factor, based on research they are currently undertaking. These latter views completely supersede the estimates in our case study, as they are based on much more detailed case studies - including on the same funding opportunities - and further analysis by GiveWell. Our case study is shared for purely illustrative purposes, as a concrete example of applying our model when aiming to benefit the current generation.

2.6.1. Case study: investing 10 years ago and giving now

Would it have retrospectively been better to invest \$1 million in 2010 and give that money plus returns to the highest-impact opportunities available now, or to give \$1 million to what we believed to be the best opportunities available back then? We consider GiveWell's research and recommendations in 2010 compared to 2020 to get a more concrete sense of this.

In 2010, GiveWell recommended six charities internationally,⁴⁴ one of which they still recommend: the Against Malaria Foundation (AMF).⁴⁵ The next year, they classified just 2 charities as their top ones, Schistosomiasis Control Initiative and AMF,⁴⁶ both of which they currently still recommend. Since then, GiveWell have extended their pool of top charities to a total of 8,⁴⁷ and have never again permanently removed a charity from their list, with the exception of Evidence Action's No Lean Season, which they recommended in 2017 but removed from their list in 2018 due to new evidence.⁴⁸

If the cost-effectiveness of the work that these top charities have been doing hasn't changed substantially since 2010, this suggests that the difference in cost-effectiveness factor has been close to 1 over this ten-year time period: if you had to give to the highest-impact funding opportunity available in 2010 (and more so in 2011) you would expect nearly the same impact per dollar as you expect today.⁴⁹

GiveWell's first published cost-effectiveness estimates⁵⁰ that are somewhat comparable to their current estimates date back to 2012. Naively, we can use these numbers to calculate how many lives(-equivalent) one could save in 2010 and in 2020. We should note, however, that GiveWell's methodology has been updated substantially over the years. Furthermore, GiveWell uses moral weights⁵¹ to arrive at estimates of their composite measure of cost-per-death-of-an-under-5-year-old-averted-equivalent, which aims to, in addition to life-saving health benefits, include other improvements to people's lives.

To minimize variance and bias due to changing methodology and moral weights over time, we look at the direct cost-per-death-of-an-under-5-year-old-averted rather than the cost-per-death-of-an-under-5-year-old-averted-equivalent.⁵² For AMF, GiveWell estimated the cost-per-death-of-an-under-5-year-old-averted to be on the order of \$2300 in 2012⁵³ and on the order of \$3500 as of the end of 2019.⁵⁴ This is equivalent to an eight-year difference in cost-effectiveness factor of roughly 0.7, and by naive extrapolation a ten-year factor of roughly 0.6.

On the financial returns side, we see that this was an exceptionally profitable decade: \$1 million invested in the S&P 500 in 2010 would have grown to roughly \$3.6 million in 2020.⁵⁵

Using these naive estimates yields that, in expectation, giving \$1 million in 2010 would have saved the lives of almost 500 children under five, whereas investing the \$1 million in 2010 and giving now would have saved the lives of close to 1000 children.

This does not yet account for value drift or risks of loss of ownership. However, we do not expect these considerations to meaningfully differ from those in the long-term impact case, in which we estimated the persistence rate to be 0.91 on average when making use of a Fund. So investing the \$1 million in 2010 and giving now could - in expectation - have saved the lives of close to 900 children.

The example hence suggests that investing in 2010 and giving in 2020 would, in hindsight,⁵⁶ have been a higher-impact strategy than giving in 2010,⁵⁷ if your aim was to

directly benefit the current generation. However, given the naive estimates we have used, we see this as an illustration rather than as a strong piece of evidence. We defer to GiveWell's own research and all-things-considered views to estimate the difference in cost-effectiveness factor for 2020 to 2030.

2.6.2. GiveWell's views

We reached out to GiveWell with the following question, which can be seen as a highly condensed version of our survey for the long-term impact case:

If 'x' is the impact of your last dollar granted in 2020, then what do you expect the impact of your last dollar granted in 2030 to be?

GiveWell responded that they have recently been conducting research directly relevant to this question, and shared a summary of their findings with us.⁵⁸ Based on their current analysis of cost-effectiveness at different levels of funding and a series of historical case studies from 2000 to 2020, they have made best-guess projections of the cost-effectiveness of their last dollar in future years, given a variety of total amounts of money that they could move to their recommended funding opportunities.

They estimate the cost-effectiveness of their last dollar in 2020 to be ~10x the cost-effectiveness of direct cash transfers in 2020. For 2030, assuming GiveWell's total money moved to their recommendations in that year will be between \$100 million and \$1 billion,⁵⁹ which is their best guess, they estimate the cost-effectiveness of their last dollar - in real terms, indexed on 2020 dollars - to be between 5x and 9x that of direct cash transfers in 2020. In nominal terms, using a 2% projected yearly inflation rate,⁶⁰ the estimated cost-effectiveness of their last dollar is between 4.1x and 7.4x that of direct

cash transfers in 2020. This suggests a median estimate for a difference in cost-effectiveness factor of ~0.6 over the next ten years.

2.6.3. Impact ratio in our model

GiveWell's projections suggest that the difference in cost-effectiveness factor will be lower from a current-generation-focused perspective than from a long-term perspective.

This difference is probably best explained by a lower historical and projected exogenous learning rate. Research into funding opportunities in this space is easier to do well, has already received a lot more time and resources, and is supported by a wealth of intervention research from development economics, which has been around even longer.⁶¹ We would hence, on the margin, expect to learn less in the coming ten years than we would in long-term impact focused funding opportunities research, which in many ways is still in its infancy.

Based on GiveWell's view, our median estimate for the difference in cost-effectiveness factor to be factor is ~0.6. We estimate the expected difference in cost-effectiveness factor to be higher, roughly at 1.0, because of the gain-loss asymmetry discussed in Section 2.5: in terms of their contribution to the expected difference in cost-effectiveness factor, the scenarios with a larger difference in cost-effectiveness factor than the median⁶² will weigh more heavily than those with a lower factor. Our exact upwards adjustment from the median is highly subjective and uncertain though.

Combining these estimates with those for the financial returns and persistence factors yields a median ten-year impact ratio estimate of 1.9 * 0.92 * 0.6 = 1.0 and an expected ten-year impact ratio of 2.3 * 0.91 * 1.0 = 2.1 in our model. This suggests that investing to

give could also be promising from the perspective of benefitting the current generation pending the further considerations discussed in Chapter 3 - but less so than when trying to benefit the long term.

2.7. Averting animal suffering

In this section, we examine the potential of investing to give in ten years from the perspective of averting animal suffering in the near term.

Here, we have less direct evidence on the cost-effectiveness of various funding opportunities and how they have changed over time than in the case of benefitting the current generation, but more so than in the case of benefitting the long term.

To get a rough idea of the difference in cost-effectiveness factor, we reached out to two impact-focused funding opportunity evaluators in the space: Farmed Animal Funders⁶³ and Animal Charity Evaluators.⁶⁴ Similar to GiveWell, we asked them the following question:

If 'x' is the impact of your last dollar granted in 2020, then what do you expect the impact of your last dollar granted in 2030 to be?

We asked them whether they could provide both a median estimate and a 90% credible interval.

2.7.1. Farmed Animal Funders' views

On Farmed Animal Funders' side, we received a collective response from their staff Kieran Greig and Mikaela Saccoccio.⁶⁵ Their estimate, which they emphasize was very quick and tentative, is a median of 0.5x and a 90% credible interval of [.17x, 2.25x]. However, from

their explanation and the toy model that they used to inform their answer, it is clear that they think this is mostly due to high compound returns over time on investment-like funding opportunities, though they might consider a broader range of funding opportunities to be investment-like than we do.⁶⁶ For example, they note:

"the potential for changing the movement's trajectory, research leading to improved allocations, fundraising leading to increased money in the movement (>\$1.10 per \$ invested), organizations report[ing] more than a 10%/year discount rate on future donations, and the growth rate of the movement itself (>10%/year), all seem to indicate to us that compound impact seems to grow faster than compound interest."

Moreover, in their toy model, their best guess estimate for the yearly exogenous learning rate (2.5%) is higher than their best guess of the rate of decreased availability of opportunities due to diminishing marginal returns on the community level (1.75%).

On learning, both endogenous and exogenous, they note:

"For animals, we think being better informed in the future is a strong consideration. Though we think it might more suggest donating 5 years from now, or funding research, after which we hope we will be much better informed. [...] We think much of the uncertainty around the impact of groups could be resolvable after 5 years."

All of this suggests that they would estimate a larger difference in cost-effectiveness factor if investment-like giving opportunities were excluded. It seems plausible that if not

the median, then at least their average estimate for the factor would be greater than 1, at least for the coming five years.

However, overall, including investment-like funding opportunities – which they think to be a significant part of the current funding opportunities available – they think investing to give is currently worse than giving now on the margin:

"Our impression for most of the major givers in this space is that they seem to mainly save in order to give later (e.g., ~5% or less allocation to giving each year). We are not sure what the overall split should be, but it does feel like the pendulum across funders is probably too far in the direction of saving right now."

2.7.2. Animal Charity Evaluators' views

On Animal Charity Evaluators' side, we received a response from their research manager Jamie Spurgeon,⁶⁷ which he emphasized was based only on a one-hour discussion with three other members of their research team. He estimated a median of 0.9x, and a 90% credible interval of [0.05x,16x].

He also listed some of the considerations they had discussed with their team, which were mostly related to exogenous learning and specific developments in the wider animal advocacy movement driving the availability of opportunities. For instance, he mentioned "improvements in technology, particularly regarding cell-cultured meat", the "proportion of charities working in higher priority countries" and the "likelihood of average intervention effectiveness increasing in different outcome categories (i.e. increased availability of animal-free products, improved welfare standards, decreased consumption of animal products, increased anti-speciesist values, and a stronger animal advocacy movement)" as relevant factors to consider.

2.7.3. Impact ratio in our model

As with the current-generation-focused perspective, the difference in cost-effectiveness factor is likely lower when your aim is to avert near-term animal suffering than when it is to benefit the long term. However, the estimates above suggest it might be higher than the difference in cost-effectiveness factor in the current-generation-focused case.

This is again largely explained by exogenous learning, as FAF's comments emphasize: the impact-focused animal welfare community is much younger than the one focused on benefitting people in the near term, it is slightly harder to do impact-focused funding opportunities research, and there is less of a foundation of existing intervention research to build on than in the case of benefitting people in the near term.

That said, as is clear from ACE's comments in particular, the availability of opportunities will likely also play an important role here, though it isn't clear whether and to what extent developments that will cause new funding opportunities to be created will outpace the effect of more funding entering the area.

Combining the estimates from ACE and FAF, and adjusting those upwards to exclude investment-like giving opportunities, we would very roughly estimate a median difference in cost-effectiveness factor of ~1.0, and an average factor of ~2.0.

In our model, this yields a median ten-year impact ratio of 1.9 * 0.92 * 1.0 = 1.7 and an expected ten-year impact ratio of 2.3 * 0.91 * 2.0 = 4.2. Hence, pending the considerations

discussed in Chapter 3, and with the important caveat that investment-like giving opportunities might be an extra strong alternative in this space, investing to give looks promising from the perspective of averting animal suffering in the near term as well.

3. Further considerations on investing to give

Here we discuss considerations beyond our proxy question that are relevant for the decision to invest to give or give now, both from a long- and near-term perspective. These are:

Further considerations that favour giving now:

- Compound returns on giving and investment-like giving opportunities
- Smaller factors excluded from our model

Further considerations that favour investing to give:

- Option value
- Higher-return investment options
- The impact of investing itself
- Others' time preference

Further considerations that increase our overall uncertainty:

- Correlations among factors
- Model uncertainty

3.1. Considerations favouring giving now

3.1.1. Compound returns on giving

A factor to consider in favour of giving now is whether there might be giving opportunities that themselves have equivalent or even larger compound returns than investing does.

An argument often mentioned in this context is that certain global health and poverty interventions may have compound social returns for beneficiaries that outweigh compound investment returns.⁶⁸ However, we should expect these returns to be bounded from above by the world economic growth rate in the short to medium term.⁶⁹ The main reason is that, even though a beneficiary or group of beneficiaries might obtain gains above the world growth rate (~3%)⁷⁰ for a few years, they will partially use these gains for consumption. More generally, beneficiaries will not indefinitely keep reallocating their gained resources (health, money, knowledge, etc.) to large-compound-return activities or financially investing those resources. Moreover, even if they were to do this for a while, they are unlikely to ultimately choose to spend their gained resources in the highest-impact ways possible to benefit others.

The same holds true for any intervention that adds some fixed amount of resources to the world and where there is no way to substantially influence how the resources that have been created are reinvested or spent.

It might hence seem that we should discount the results of our model by, at most, an expected global economic growth rate of about 3% per year,⁷¹ which amounts to a downwards adjustment for the impact of investing to give in ten years of roughly 26%.

Such an adjustment, however, has technically already been taken into account in our estimates of the difference in cost-effectiveness factor.

More generally, it is certainly possible that there are funding opportunities with effects that extend and perhaps compound far into the future: these are precisely the types of funding opportunities we are trying to identify when aiming to benefit the long term. For example, we recommend funding opportunities for existential risk reduction for their potential to do so.⁷² However, for the existence of these funding opportunities to be favouring giving now, it needs to be the case that their *total* benefits outweigh those of investing to later give to funding opportunities whose impact similarly, and perhaps even more effectively, extends into the far future. The fact that current funding opportunities can start accumulating their effects ten years *earlier* does not imply that their effects persist *longer overall* than those of funding opportunities in ten years from now. And whether these effects do persist longer overall is the type of comparison that we are already making when trying to answer our proxy question via our model. So the existence of longer-term compound returns on giving opportunities is not an *additional* argument to give now over investing to give, but is already included in our attempt to answer our proxy question.

Investment-like giving opportunities

There are, however, some giving opportunities which we labelled "investment-like" and intentionally excluded from our proxy question, because funding them is in some ways more similar to investing to give than to funding other giving opportunities. These are giving opportunities whose primary route to impact is making more financial or human resources available to be "spent" on the highest-impact opportunities broadly - not just on a particular, pre-specified problem - at a later point in time.⁷³ We think these opportunities

could in principle have higher compound returns than investment, even in the longer term, mainly because they are able to benefit from exogenous learning as much as investing to give does.

The most straightforward example of an investment-like giving opportunity is encouraging other people with aligned values to invest to give rather than not give at all: at a high enough success rate per dollar spent on this, it is clear that this would beat direct investment, and the main uncertainty is whether such a success rate can be achieved. More general effective altruism⁷⁴ movement-building may also qualify, but this is only true with certainty if it leads to a high enough rate of people joining and a high enough rate of those people making close to the highest-impact choices, for example regarding investing to give versus giving now.

Capacity building, including endogenous learning, could be another example: if certain activities - for instance global priorities research⁷⁵ - increase the amount of resources that are allocated to what are *actually* the highest-impact opportunities at later points in time and if further capacity building is able to build on these gains, those activities could have large compound gains as well. However, capacity building only offers such compound gains if it wouldn't happen exogenously in the near future anyway.⁷⁶

If we find current investment-like giving opportunities that seem to provide higher returns than investing, we should still consider the availability of such opportunities and our ability to detect them over time. That is, if we expect better investment-like giving opportunities to become available over time or expect to get better at identifying them, it might still be better to invest. However, this is not the case if current investment-like giving opportunities would themselves lead to someone else giving a large enough amount to the better investment-like giving opportunities that might become available later. For example, if some currently available investment-like giving opportunity caused an extra philanthropist with similar aims to give \$10 million in 10 years from now, this new philanthropist could spend that money on the best investment-like giving opportunities available at that time, so there would be no loss by giving to the currently available investment-like opportunity compared to investing the funds financially.

All of this said, we are generally less confident in our reasoning about compound returns on giving and investment-like giving opportunities than about the other considerations in this chapter, given the little time we have been able to spend looking into these topics relative to their complexity. We see this as an important area for further research.

3.1.2. Smaller factors excluded from our model

We have excluded some factors from the model because we think they would not meaningfully influence our overall results. However, overall we think these factors push slightly in the direction of giving now. They include taxes, investment and fund management costs, and diminishing returns to giving within our own spending.

Tax considerations

Tax considerations may be relevant for certain individuals, but there are ways, for example the use of a Donor-Advised Fund (DAF),⁷⁷ which can largely decouple them from timing-of-giving choices. A Founders Pledge Long-Term Investment Fund would likely have properties similar to those of a DAF, so that any tax benefits materialise at the time of allocation to the Fund rather than at the time of disbursement from the Fund to funding opportunities.

Administrative costs

We have not included investment and fund management costs because we believe these to be quite small compared to investment returns, certainly when pursuing index investing. We would however not be surprised if all administrative costs combined, certainly if research costs are included, amount to one or more percentage points over the course of a decade, depending on the size of a Long-Term Investment Fund.

Diminishing returns

We expect diminishing returns to our spending within a particular year to be negligible for practical purposes, because, as discussed in the Introduction, our giving should be seen as additional to the total giving of a community with aims similar to ours, and it will be dwarfed by that in most cases. This is more so when trying to benefit people currently alive than for animal welfare and long-term-focused efforts, because there are a lot more resources dedicated to the former.⁷⁶ However, it seems justified for the latter as well, as impact-focused grantmaker Open Philanthropy alone currently makes grants with the aim to maximally benefit the long term on the order of \$100 million per year.⁸⁰ It hence seems unlikely that our first dollar spent in any year would have substantially different expected impact from our last in that same year.

3.2. Considerations favouring investing to give

3.2.1. Option value

When choosing to invest to give, we don't need to commit to disbursing the funds in exactly 10 years from now. We could choose to disburse them next year, or in 5 years, 50 years or 5000 years from now. The only option we lose is to give now. However, if we choose to give now, we lose all those other options, and the opportunity to learn about which option is best in the meantime.⁸¹ Investing hence has more option value than giving now does.

Option value has multiple important implications for the potential of investing to give.

It is, first of all, important for interpreting the expected impact ratio of investing to give vs giving now in our model. The ratio provided by the model serves as a proxy for the more general question of whether investing to give is the highest-impact philanthropic strategy available, but it doesn't *necessarily* answer that question, depending on the ratio we find. If the ratio is higher than 1, this suggests that investing *is* the highest-impact strategy available at this point in time, excluding investment-like giving opportunities. However, if the ratio were lower than 1, this would *not* imply that investing is not the highest-impact strategy: it could, for instance, be that, because of changing circumstances over time, our expected impact from giving in 10 years is lower than our expected impact of giving now, but that our expected impact from giving in 100 years is higher. More generally, if we can identify just one point in time in the future such that our expected impact of investing to

give then is higher than our expected impact to give now, this suffices to make investing the optimal strategy.

A second and somewhat overlapping way in which option value is relevant is in how it allows us to use the "watch then pounce" strategy.⁸² Even if our ability to have an impact were to decrease over time *on average*, it might be that, every once in a while, unusually high-impact funding opportunities arise. For instance, there might be a limited time window in which particularly beneficial government policy can be passed or a new institution can be set up. When investing, we have the option to fund such opportunities, whereas giving now restricts us to the funding opportunities available now.

A third and final way in which option value is relevant, is in how we are able to learn more about the answer to "When do I give?" over time. If we wouldn't expect to learn anything more about which option is best, we would have to choose one option now, and even though "giving now" would seem unlikely to be the very best *a priori*, we might conclude that it is the best one given the limited information that is available. However, when investing, in addition to learning about higher-impact funding opportunities, we can expect to learn more about when is the best time to give. Hence, even if we concluded that given *current* evidence, "giving now" looks more promising than giving at any other time, the option to access future evidence that investing provides could make it a higher-impact strategy in expectation.⁸³

3.2.2. Higher-return investment options

We have conservatively assumed equity market index funds as our investment strategy. There are alternatives that might yield higher returns, such as leveraged index investing or venture capital investing.

Leveraged investing⁸⁴ in index funds is a higher-investment-risk-per-dollar and higher-expected-return strategy than regular index fund investing, but it relies on the same underlying mechanism for its returns. The higher investment risk should matter a lot less to us than to the self-interested investor, as we should see our financial resources as part of those in a community of people and organisations with similar aims. However, it still matters to some extent, for instance because the performance of our investment portfolio is likely correlated with that of others in the community.⁸⁵ Another important practical consideration is that, due to restrictions on higher-investment-risk investments by nonprofits in many countries, it might be difficult to set up a charitable investment fund that pursues this strategy for a large part of its investments.

Venture capital investing is also riskier than regular index investing, and the same practical consideration on investment risk restrictions for charitable funds holds. However, reliably achieving higher expected returns is probably only possible with a specialized skill set, access to the right markets, and access to the right information. On the other hand, if one has some or all of these, the expected returns for venture capital investing could arguably be much higher than for index investing, even taking into account increased investment management costs.

We will consider both these alternatives for future research and in our plans for future investment strategies of a potential Long-Term Investment Fund.

3.2.3. The impact of investing itself

In addition to the impact from giving, there could be an impact from the investment itself. The additional impact of index investing is probably negligible for our purposes, but this could be different if one takes an impact investment strategy.

We have previously published a report on the potential for impact of impact investing,⁸⁶ in which we make a comparison with giving now. We conclude that having a larger impact via impact investing than giving now is much more challenging than it may seem, and that we have yet to find examples of reliably higher-impact impact investments than our recommended giving opportunities. However, we see no reasons for those not to exist in principle, and the report provides considerations for finding the highest-impact impact investments.

Combining those conclusions with the findings in this report suggests that a hybrid impact investing and investing to give strategy is at least worth exploring. As with leveraged and venture capital investing, this is something that we will consider for future research and in our plans for future investment strategies of a potential Long-Term Investment Fund.

3.2.4. Others' time preference

Lastly, there is a wholly separate, theoretical argument for investing to give, rooted in the hypothesis that many people have an intrinsic preference to receive benefits now over receiving those in the future. To what extent this is true is a matter of substantial academic

debate,⁸⁷ but most economists seem to agree that there is very likely some positive rate at which most people intrinsically discount future benefits.⁸⁸

Time discounting by others makes it more plausible that investing to give is a high-impact strategy in two distinct ways. First, it helps explain why we are able to obtain large returns on financial investments. If most people value the future less than us, they will be willing to trade their influence over the future for more influence over the present, and they will not take some of the investment opportunities that are advantageous from the perspective of benefitting the long term, leaving those to us. Secondly, it makes it more likely that humanity as a whole - from the perspective of someone who cares equally about others regardless of when they live - is overspending on the present, and that we can move humanity's overall portfolio closer to optimal by investing.⁸⁹

3.3. Considerations increasing uncertainty

3.3.1. Correlations among factors

We have modelled the three key factors in our model as independent of each other, whereas in reality, there are certainly correlations between them. Positive correlations between the factors strengthen the case for investing to give, whereas negative correlations weaken the case. This is mainly because with positive correlations, the more extreme positive scenarios become more likely, and those will disproportionately influence the *expected* impact of investing to give. Below we run through all potential two-way correlations in turn. Overall, we don't see a clear case for bias in a particular direction by leaving out these correlations, but they do substantially increase our overall uncertainty.⁹⁰

Financial returns and persistence

Financial returns and persistence of funds and values are likely positively correlated, for instance because better financial returns could on balance add to the motivation of an individual or organisation to follow through on their intentions. However, we don't expect this to be a major effect.⁹¹

Financial returns and difference in cost-effectiveness

We expect a small positive correlation between financial returns and the exogenous learning component of the difference in cost-effectiveness factor, as it seems likely that investment returns will be somewhat positively correlated with how many resources are being spent on funding-opportunity-relevant research globally.

However, we are highly uncertain about the direction and strength of the correlation between financial returns and the availability of opportunities.

On the one hand, higher financial returns will in the longer term be correlated with more funding being available in the community of people and organisations with aims similar to ours. Given diminishing marginal returns to spending at a community level, this could significantly reduce the availability of high-impact opportunities.⁹²

On the other hand, higher investment returns will be correlated with faster economic growth and technological progress, which could provide new opportunities to have an

impact, not only by making it more likely that new high-impact projects will be initiated more generally, but also by creating more risks to address and problems to solve, and in turn bringing about new ways to solve these problems. Climate change is a pertinent example: economic growth and technological progress have played a major role in causing this problem,⁹³ but also have the potential to play a major role in driving solutions.⁹⁴

Furthermore, the amount of funding that is available also has a (partially compensating) positive effect on the availability of opportunities: a better funding climate for nonprofits would allow more high-impact funding opportunities to start, in the standard way that more demand for a good, in addition to driving up its price, can lead to an increase in its supply.

Lastly, the timing of extra spending by other philanthropists in the community matters: do they think carefully about when to spend the extra money themselves? For example, if they spend all their gains shortly after they occur, this would only temporarily reduce the availability of opportunities. We could then wait and spend our funds at a time when there is less overall funding available.

Overall, we think financial returns and differences in cost-effectiveness are meaningfully correlated. We have high uncertainty about both the sign of the correlation and its size, but think a negative correlation is slightly more likely.

Persistence and difference in cost-effectiveness

Persistence and exogenous learning are unlikely to be related strongly. The only potential interactions seem to point to a slightly positive correlation: more exogenous learning could

be motivating, and better persistence of funds and values could drive a motivation to learn more.

A possible interaction between persistence and availability of opportunities is for there to be both more loss of ownership risks and more high-impact giving opportunities in times of high risks of catastrophe. We didn't identify any other clear ways in which the two are related, and this way accounts for only a small subset of scenarios, so we expect at most a small negative correlation among the two on balance.

Overall, then, we don't expect there to be a strong positive or negative correlation between the persistence and difference in cost-effectiveness factors.

3.3.2. Model uncertainty

Lastly, it is worth emphasizing that our way of approaching this complex question generally comes with a lot of uncertainty and room for error, both with regards to the individual factor estimates and the overall model structure and assumptions, and in ways that we might not be even aware of. Our conclusion here is hence necessarily tentative, and we hope to keep improving our view on investing to give through future research.

4. Conclusions and recommendations

We find a promising case for investing to give as a high-impact philanthropic strategy. This is true across all three aims explored in this report, but most so when aiming to benefit the long term.

4.1. Benefitting the long term

In Chapter 2, we estimated a median impact ratio of 2 and expected impact ratio of 9 for our proxy question of investing to give in ten years compared to giving now. A lot of these expected benefits are driven by exogenous learning: our increased ability to detect high-impact opportunities over time. Financial returns are the second-most important driver.

In our sensitivity analysis, we discussed that there does exist a somewhat plausible view you could hold to arrive at an *expected* impact ratio estimate lower than 1 - based on a heavily decreasing marginal cost-effectiveness over time - but that under many other plausible views investing to give in ten years comes out as higher-impact in expectation.

As explained in the Introduction and in Section 3.2.1 on option value, these estimates for investing to give in ten years likely underestimate the expected impact of investing to give more generally: one doesn't have to fix the timing of your disbursement at ten years from now. Furthermore, it might be possible to make even larger investment returns than via equity index investing and/or increase one's impact via impact investing, though the latter is more uncertain. And there is a separate argument for investing to give, based on others' time preference.

On the other hand, even though compound returns from most giving opportunities do not seem to meaningfully affect the case for investing to give, compound returns from *investment-like* giving opportunities weigh in favor of giving now. Investment-like giving opportunities with longer-term returns that outperform financial returns seem to be the best competitor to investing to give.

Furthermore, our report has approached investing to give from the perspective of an individual philanthropist, and has assumed that their investing or giving decisions do not meaningfully change the investing vs giving distribution at the level of the community of people and organisations with similar aims. This was a crucial assumption for us to be able to make practically relevant estimates, but also limits the strength and durability of our conclusion: if, for instance, this report itself causes a lot more philanthropists to invest to give, this could change our conclusions.

All of this said, we should emphasize the uncertainty that comes with these estimates and claims: this research project is among the first investigations of investing to give that we know of, and we wouldn't be surprised if we have made substantial errors or will change our views with further insights. We are particularly uncertain about our estimates of the difference in cost-effectiveness factor, given the limitations of our survey, and to what extent returns to giving now can compound over time, given the complexity of the topic.

We hence think the right conclusion for us to draw is that investing to give is a very promising strategy for any long-term oriented individual philanthropist to consider at this point in time. We do not think we can confidently claim it is the highest-impact philanthropic strategy, but we think it is plausible that it could be, and we recommend our members to at least consider making it a part of their philanthropic portfolio.

For Founders Pledge in particular, the conclusions of this research project are a strong argument to set up a Long-Term Investment Fund for our members which is optimized for financial returns and persistence, and which will be able to make use of new insights on high-impact funding opportunities and investing to give itself as they come in over the years. We have already started exploring the practicalities of doing this, and cautiously expect to launch such a Fund in 2021.

4.2. Benefitting the near term

In Chapter 1, we estimated median ten-year impact ratios of 1.0 and 1.7 and expected ten-year impact ratios of 2.1 and 4.2 when aiming to benefit the current generation and aiming to avert animal suffering respectively.

Most of the further considerations that apply when aiming to benefit the long term also apply here, including those of option value and compound returns on investment-like giving opportunities.

We think that for benefitting people in the near term, our assumption of not meaningfully changing the distribution of investing vs giving at a community level is even more justifiable than in the long-term impact case, as the relevant community is larger.⁹⁵ For the same reason, we think it less likely that we are able to find great investment-like giving opportunities in this space. On the other hand, our estimates for the median and expected ten-year impact ratios are considerably lower, though still larger than or equal to 1. Given all this, we think investing to give is an option worth considering for our members aiming to benefit the current generation, but less so than for our members who aim to benefit the long term, and only when a suitably designed investment and granting vehicle becomes available.

For averting animal suffering in the near term, our assumption of not meaningfully changing the distribution of investing vs giving at a community level seems less justifiable: we don't know of other impact-focused grantmakers than Open Philanthropy who put at least a few tens of millions of dollars into animal welfare each year.⁹⁶ And, as Farmed Animal Funders noted,⁹⁷ there seems to be a lot of promise for investment-like funding opportunities in the space. On the other hand, our median and average estimates for the ten-year impact ratio are higher than in the case of benefitting people in the near term. Taken together, we think investing to give when focusing on near-term animal welfare is also an option worth considering, and think it currently looks more promising than giving to non-investment-like giving opportunities. However, we expect there to be investment-like giving opportunities that can outperform or at least match investing to give, and are generally less confident in our conclusions here than in the case of benefitting the long term or the current generation of people.

4.3. Suggestions for future research

As we have emphasized, research into this topic is still in its infancy, and we expect our views to change over time. Here we highlight a few directions of future research that we think would be especially likely to accomplish such change. We plan to pursue some of these directions ourselves, but would be very happy to see others do this as well.

4.3.1. Critiquing our assumptions and methodology

We strongly invite any external criticism, as this is a complex issue and we wouldn't be surprised if we have made important mistakes in our treatment of it.

4.3.2. Improving our difference in cost-effectiveness estimates

The difference in cost-effectiveness factor is a crucial part of our analysis, we are still very uncertain about our estimates of it, and there are some clear ways to improve those estimates. In particular, running an improved version of our survey - accounting for the limitations discussed in Appendix IV - on a larger and more diverse sample of experts seems valuable. In addition, there are other ways to make estimates, such as studying historical trends, which could be a valuable complement to the consultation of experts.

4.3.3. Updating our estimates over time

We think it will be important to update our estimates at least every few years or so, not only to account for new learnings but to account for changing circumstances as well. As discussed above, our estimates for each of the factors could change significantly over time, for instance if financial market circumstances change, if the exogenous learning rate slows down, or if a lot more (or less) philanthropists start to invest to give.

4.3.4. Explicit modelling of the future

Our estimates in this report are mostly informed by either extrapolation of historical trends or direct consultation with experts. We think those are enough to provide valuable initial insights, but that they could be complemented with explicit models of the world that examine how different empirical assumptions and scenarios would affect each of the relevant variables. Making such models could provide generally more accurate estimates, but perhaps more importantly could improve our understanding and unearth any faulty assumptions we are currently making in our thinking about this topic. Furthermore, more complex and explicit models could shed more light on how important correlations among the different key factors are.

4.3.5. Modelling at the community level

As explained in the Introduction and as emphasized in our conclusion, two of the most important choices we have made in this project were to approach the question from the individual rather than community perspective, and to assume that individual spending will not meaningfully affect the community distribution between investing and giving now. We think alternative approaches, such as the one Philip Trammell takes in his paper,⁹⁸ hold a lot of promise. They face their own challenges in having to make other assumptions, but could at the very least complement the approach presented in this report and compensate for some of the biases our approach will undoubtedly have.

4.3.6. Researching compound returns on giving

We have discussed some important considerations on compound returns on giving and investment-like funding opportunities in Section 3.1.1., but we think this warrants a lot more research, not in the least because we conclude that funding investment-like funding opportunities may be the best alternative to investing to give. This is an intricate topic, and we would not be surprised if our current view of how benefits do or do not compound over time is mistaken in important ways.

4.3.7. Investigating alternative investment strategies

Throughout this report, we have assumed index investing as our investment strategy, and we have based our estimates for investment returns on historical index fund performance. As discussed in Sections 3.2.2 and 3.2.3, we think there are alternative strategies that might yield more financial returns or impact in expectation, which we will consider in our plans for a Founders Pledge Long-Term Investment Fund. Hence, research into the potential of these strategies in the context of investing to give would be valuable.

Appendix I: Financial returns estimate

Historical return rates

To start narrowing down our estimate, consider that the S&P 500 index has had an inflation-adjusted annualized⁹⁹ return of ~7% since its inception in 1926.¹⁰⁰ We need to adjust this estimate for selection and survivorship bias, as there have been multiple markets in other countries that have done a lot worse, or have even ceased to exist, such the Rio de Janeiro Stock Exchange.¹⁰¹ A recent Credit Suisse report adjusts for this and considers an index of all equity globally. It finds an annualized real return on global equity of 5.2% from 1900 to 2019.¹⁰²

Uncertainty

To determine the appropriate level of uncertainty in our estimate, we consider the annualized real returns for the S&P 500 for the nine 10-year periods from 1 January 1930 to 31 December 2019 and naively account for selection and survivorship bias by subtracting 1.8%:
| Period | Annualized real returns (<u>source</u>) | Annualized real returns, adjusted for bias | Ten-year real returns factor, adjusted for bias |
|-----------|--|--|---|
| 1930-1939 | 2.0% | 0.2% | 1.0 |
| 1940-1949 | 3.5% | 1.7% | 1.2 |
| 1950-1959 | 17.0% | 15.2% | 4.1 |
| 1960-1969 | 5.1% | 3.3% | 1.4 |
| 1970-1979 | -1.5% | -3.3% | 0.7 |
| 1980-1989 | 12.0% | 10.2% | 2.6 |
| 1990-1999 | 14.9% | 13.1% | 3.4 |
| 2000-2009 | -3.4% | -5.2% | 0.6 |
| 2010-2019 | 11.5% | 9.7% | 2.5 |

Though the sample is too small to perform any formal statistical tests, the ten-year real returns factor data appear to very roughly fit a log-normal distribution with a median of 1.6 and a 90% credible interval of [0.60,4.1].¹⁰³

We should check these uncertainty bounds for selection and survivorship bias. It could be, for instance, that the world index has larger variance because the American stock exchange did not completely collapse in the time period under consideration, whereas some other markets did. However, the standard deviation of the yearly returns for the world index from 1900 to 2020 (0.174)¹⁰⁴ was slightly lower than the standard deviation for the S&P 500 from 1927 to 2020 (0.201),¹⁰⁵ so we don't see the need to make a substantial adjustment.

Inflation

So far, we have used inflation-adjusted returns. These are appropriate in most circumstances, because they are meant to reflect the extent to which funds have truly increased in value. However, there is some controversy about whether inflation-adjusted prices using the Consumer Price Index - which is the standard way to adjust for inflation - are accurate: it is possible that the Consumer Price Index is biased upwards, and hence that average recent global real returns have been (a lot) higher.¹⁰⁶

Moreover, for the purpose of this project we are not interested in what Mary's funds can buy in terms of consumer goods, but in terms of impact, for which the price index might be substantially different. Indeed, this 'impact price index' is what the difference in cost-effectiveness factor is meant to reflect in our model. In making our estimates for that factor,¹⁰⁷ we have chosen not to include inflation: we asked the experts that filled out our survey to provide marginal cost-effectiveness estimates with the cost expressed in nominal dollars for each year in question. We should hence reintroduce¹⁰⁸ inflation in our financial returns estimate.

The US Federal Reserve Bank commonly target an inflation rate of 2% per year,¹⁰⁹ and in the US, the inflation rate has been close to this target since the early 1990s.¹¹⁰ As we would expect a large part of Mary's funds to be invested in the US and similar markets, and with global inflation also having been close to 2% in the past six years,¹¹¹ we will use this as our estimate for the future inflation rate. This corresponds to an adjustment of about 22% over ten years.

Overall estimate

Adjusting our previous estimates for the ten-year real returns factor upwards by 22% to reintroduce inflation yields a 90% credible interval of roughly [0.73,5.0], a median estimate of 1.9 and an average estimate of 2.3 for the ten-year financial returns factor.

Appendix II: Persistence estimate

Value drift

We don't have any direct evidence on Mary's risk of value drift, but we do have data on people in similar situations and with similar aims that we can use to arrive at a ballpark estimate.

Ben Todd of career advice organisation 80,000 Hours recently reviewed various pieces of evidence on drop-out and decreased involvement in the effective altruism movement,¹¹² which shares a lot of values with Mary and Founders Pledge. He estimates a five-year decreased involvement rate of 10-40% depending on the person in question's level of engagement with the effective altruism movement, with those most engaged most likely to stay involved.¹¹³ He argues that the value drift rate for a particular person likely decreases over time and in years six to ten could be roughly 50-75% as large as the value drift rate during the first five years of involvement, which would imply ten-year value drift rate estimates of roughly 15-55%.¹¹⁴

These rates are a good starting point, but we need to adjust them to account for our specific situation.

First and most importantly, Mary can make use of a legal vehicle to improve the chances of her funds being spent on the highest-impact opportunities in 10 years from now or beyond. At Founders Pledge we already operate a Donor-Advised Fund that she could use to ensure the money will be spent on charitable efforts, and a potential Long-Term Investment Fund would be organized in a way that is optimized for long-term impact.¹¹⁵ If Mary commits her money to this Fund, which will allocate it to high-impact funding opportunities on her behalf at a later point, the question changes from what her risk of value drift is to what the Fund's risk of value drift is, and this Fund would be set up in a way to minimize that risk.

Second, 'decreased involvement' in Todd's estimates is only loosely defined. It includes cases in which someone would largely stick to their previous commitments to give to the highest-impact opportunities available, but might, for instance, give a little bit less, dedicate a smaller part of their career directly to making the world a better place, or engage less directly with the effective altruism movement while still aiming to have a maximum positive impact on the world in their own lives.¹¹⁶

Founders Pledge's primary purpose is to have a maximum positive impact on the world via supporting its members and is made up of multiple people and checks and balances which help maximise the probability of it staying on course (assuming it doesn't succumb to other risks, some of which are discussed below). The Long-Term Investment Fund would be set up in a way that is even more secure. For instance, we could consider setting up a separate legal entity and detach it from Founders Pledge's main organisation in such a way that it doesn't face many of the challenges and pressures that any non-profit faces, such as fundraising. We could also consider setting up a governance structure for the Fund that further minimizes the risk of individual or group value drift leading to the Fund's value drift, by distributing decision-making power over multiple value-aligned but independent individuals or organisations both internal and external to Founders Pledge. Given all this, we think such a Fund should quite easily be able to outperform the 15% ten-year value drift rate Todd estimates for the average most engaged individual in the effective altruism community.

Loss of ownership

Loss of ownership could occur in multiple ways. A first category to consider is that of catastrophe: either personal, local or global. A Long-Term Investment Fund protects against any personal risk to Mary, such as theft or death. However, it could lose funds if the legal system breaks down and/or property rights are no longer protected.

This would not happen in most local catastrophes, such as natural disasters and wars, except for the most extreme ones. It would however happen in some extreme global catastrophic¹¹⁷ or even existential risks¹¹⁸. In his book *The Precipice¹¹⁹*, Toby Ord surveys the latter category and estimates a risk of one out of six in the coming century, which converts into a ten-yearly risk of roughly 2%.

On the one hand, from Ord's argumentation it appears he thinks the risk is increasing over the century (most of the risk comes from future technologies), so we should perhaps revise this slightly downwards for the upcoming decade.

On the other hand, we should adjust this estimate upwards for all non-existential local or global catastrophic risks that would be sufficient for the funds to be lost. To put a rough upper bound on the magnitude of those risks, we can use historical data on the breakdown of states and empires. The breakdown of a state is not a necessary condition for the breakdown of a legal system and/or violation of property rights, but it is further from a sufficient condition: there are enough cases in which the breakdown of a state wouldn't lead to loss of ownership that the order of magnitude of state breakdown should serve as an appropriate upper bound for risk of loss of ownership. In his forthcoming book, Anders Sandberg analyzes the persistence of empires in the past 5000 years and European states

between the years 1000 and 1850, and finds roughly exponential fits with half-lives in the range of 100 to 400 years,¹²⁰ which translates into 10-year breakdown rates of about 3 to 10%. Combining this upper bound with the lower bound of the existential risk estimate, we expect the risk of loss of ownership from catastrophe to be in the order of a few percentage points per ten years.

A second category is that of government policy and taxation changes. These changes seem more likely to occur than relevant catastrophes. For instance, recently minimum payouts for charitable trusts were <u>discussed</u> at the highest levels of the UK government. However, the effect of changes like these on the question of whether to invest or give now is very limited. This is firstly because in most cases, there would be time for the Fund to potentially change strategy. For instance, if it were announced that some relevant tax rate would be significantly increased, and this would make it better from an impact perspective to give now rather than to invest, the Fund could disburse funds before the new tax rates are implemented. Secondly, because most of the potential policy changes would only lead to either a voluntary or enforced (e.g. in case of a minimum payout) change of strategy from investing to giving now, they cannot affect the question which of the two has more of an impact by very much. Hence, we think the risks of government policy and taxation change are negligible for the purpose of this research project.

Thirdly, historically there have been multiple examples¹²¹ of outside legal challenges to funds that were set up for long-term investment. However, zooming in on these examples reveals that they were challenged on grounds that don't apply to the Founders Pledge Fund, at least in the short to medium term, e.g. personal political gain and the risk of a fund growing so large as to be a danger for the overall economy. In addition, longer-term investing is currently legally well-established and commonplace, at least in some countries. For example, charities in the UK hold over 140 billion pounds in long-term investments,¹²² and just one charity, the Wellcome Trust, has more than 30 billion pounds of assets currently invested.¹²³ It hence seems very unlikely that a Long-Term Investment Fund would face outside legal challenges like these, at least in the first decade of its existence. And even if such a challenge would arise, a similar argument applies as for government policy changes: often, challenges could be anticipated and there is the option to (partially) change strategy to giving now. Overall, we therefore conclude that this risk is negligible as well, at least for the first ten years of existence of a Fund.

Lastly, there is the risk of Founders Pledge itself no longer existing in 10 years from now, for reasons other than the ones already accounted for. At first, this might seem like a significant risk, given for instance an average yearly exit rate of nonprofits on the order of 3 to 5% in the US,¹²⁴ equivalent to 25 to 40% in ten years. However, the Long-Term Investment Fund can be set up as a separate entity, for example as a <u>Charity Authorised</u> <u>Investment Fund</u> in the UK, and risks for the funds in the Fund can be largely decoupled from risks for Founders Pledge. In addition, here again, the considerations of being able to anticipate these risks and being able to change strategy hold. We hence judge this particular category of risk of loss of ownership to be no larger than a percentage point over ten years.

Overall estimate

As the different reasons leading to value drift and loss of ownership appear to be largely independent, we can arrive at a final estimate of the persistence factor by multiplying together the persistence rate for each of those reasons. In our short analysis of the reasons above, we have found that value drift and loss of ownership due to catastrophe – both on the order of a few percentage points – are the dominant reasons. Combining those and accounting for the considerable uncertainty in our rough analysis (including the existence of "unknown unknowns"), we estimate a 90% credible interval for the persistence factor of [0.8, 0.98] and a median of 0.92. Fitting this to a beta distribution¹²⁵ yields an average persistence factor of 0.91.

Appendix III: Survey questions

Please see below for the full outline of the survey titled "The cost-effectiveness of marginal funding opportunities over time", which was administered using Google Forms.

Section 1: The cost-effectiveness of marginal funding opportunities over time

This survey is part of a Founders Pledge research project on the value of investing as a philanthropic strategy. You can look at this post

(https://forum.effectivealtruism.org/posts/CfLoq8nJBzRARohtQ/the-case-for-investingto-give-later) for an idea of what this research project entails, but please only do this after having filled in the survey, to avoid being anchored to particular estimates. Your response to this survey will be an important input to this research project. We are able to process any responses that are submitted by the 24th of September.

We appreciate you taking the time to fill this in. The survey should take about 45 minutes to complete. It consists of three sections, with the second containing the main questions.

Your email address and name are recorded only to be able to verify your response, and will not be stored or used for any other purpose unless you indicate at the end of the survey

- that you are happy to be contacted to answer any follow-up questions

- that you would like to receive the final report once it's finished

- that you are ok for your name to be mentioned in the acknowledgements for the final report

In particular, for the purposes of the project and final report, your name and email address will not be linked to the remainder of your response in any way, and your background only insofar it links your responses to a relevant category, e.g. 'charity researchers'.

Elaborate instructions for making estimates will be provided in the next section.

For any questions or comments, please reach out at sjir@founderspledge.com.

Email address

[answer]

Name

[answer]

Please provide a few bullet points on your background and expertise in evaluating (the cost-effectiveness of) philanthropic funding opportunities, including the causes you have relevant expertise in and for how long you have been working on funding opportunity evaluation.

[answer]

Which of these worldviews currently has a dominant influence on your personal judgment of the expected cost-effectiveness of funding opportunities?

Please use the rough definitions listed in this post:

https://www.openphilanthropy.org/blog/update-cause-prioritization-open-philanthropy, and answer the questions from your personal perspective, even if your main area of expertise is not (directly) relevant to that. E.g. if your expertise is mainly in evaluating animal welfare funding opportunities, but you think the most cost-effective funding opportunities are to be found using a long-termist lense, please indicate 'long-termist' here.

[long-termist OR

near-termist, human-centric OR

near-termist, animal-inclusive OR

other]

Section 2: Your estimates

We are interested in your personal estimates of the cost-effectiveness of the 'marginal funding opportunities' in certain years in the past and future, relative to the cost-effectiveness of the marginal funding opportunity for 2020. Please use the below definitions and instructions to make your estimates:

EXCLUDING INVESTMENT-LIKE FUNDING OPPORTUNITIES

For the purpose of this survey, please exclude any funding opportunities that are 'investment-like' from your consideration. These are funding opportunities whose primary expected way of having an impact is via making more financial or human resources available to be allocated to the most cost-effective opportunities broadly (not just to a particular, pre-specified problem) at least 10 years later, e.g. effective altruism movement building and cause prioritization research.

MARGINAL FUNDING OPPORTUNITY

The 'marginal funding opportunity' for a given year is the hypothetical funding opportunity that you would, at the time, choose to spend your next dollar on after already having spent everything you are spending on funding opportunities in that year. More precisely, it is characterized by the following:

(1) It is available to you in the given year (both detectable and fundable)

(2) It isn't funded by anyone else in the given year

(3) If in the given year, you were to judge the cost-effectiveness of this funding opportunity, you would at that point conclude that it is in expectation the most cost-effective funding opportunity among those that aren't funded in that year

COST-EFFECTIVENESS ESTIMATES

The cost-effectiveness of a funding opportunity refers to how much value you think is created - according to your overall moral framework and sense of what value is - per dollar given to that funding opportunity, compared to the counterfactual of the funding opportunity not receiving that dollar.

When asked for your estimate below, please provide your best guess of the 'actual' cost-effectiveness of a funding opportunity, i.e. what you would say if you were given idealized, full knowledge of how the world works and what the 'right' moral framework and theory of value is. Please do this no matter whether the funding opportunity in question is in the past or in the future. However, please do not take into account inflation, i.e. always treat 1 dollar as 1 dollar in the year the question refers to.

Importantly, please note the distinction between the cost-effectiveness estimates of marginal funding opportunities that we are asking you to provide, i.e. "idealized" estimates, and the way we have defined what the marginal funding opportunity is in a given year, which is based on the cost-effectiveness estimate you would make in that year, based on your knowledge and values at the time. What your marginal funding opportunities in past years were depends on your personal history, and what those will be in future years depends on how you expect your knowledge and values to develop (and those of the people and organizations that influence your judgement).

RELATIVE COST-EFFECTIVENESS

We ask you to provide your estimates in relative terms to the actual cost-effectiveness of the marginal funding opportunity in 2020. Whatever it may be in reality, here that value is defined and fixed as '1'. Hence, if you think the cost-effectiveness of a marginal funding opportunity in a particular year is 'y' times that of the one in 2020 (for some specified number 'y'), please fill in 'y' for your estimate.

MEDIAN AND CONFIDENCE INTERVALS

We ask for your median estimates as well as your 90% confidence intervals:

- Your median estimate should be a number that you expect to be too high with 50% probability, and too low with 50% probability

- The lower bound of your 90% confidence interval should be a number that you expect to be too low with 95% probability, and too high with 5% probability

- The upper bound of your 90% confidence interval should be a number that you expect to be too high with 95% probability, and too low with 5% probability

PROMPTS

In addition to using the direct evidence you have from your work, these might be useful prompts to keep in mind while making these estimates:

- How do I think the general availability of cost-effective opportunities has shifted/will shift?

- How do I think my ability to find the most cost-effective opportunities has shifted/will shift, including shifts in my moral framework and values?

- How do I think the extent to which the most cost-effective funding opportunities that I can find are fully funded has shifted/will shift?

WHAT TO FOCUS ON

These estimates will necessarily be subjective and hard to explain perfectly, but we ask you to provide supporting reasoning and evidence insofar possible. Please don't spend time on making each of the individual estimates precise: the rough ballpark they are in, how they relate to each other, and your overall reasoning are much more important input than the precision/granularity of any of the estimates individually.

TEMPLATE FOR ESTIMATES (see question below)

| Median | |
|--------|--|
| 2010 x | |
| 2015 x | |
| 2018 x | |
| 2019 x | |
| 2020 1 | |
| 2021 x | |
| 2022 x | |
| 2025 x | |
| 2030 x | |

| Lower bound of 90% confidence interval |
|--|
| 2010 x |
| 2015 x |
| 2018 x |
| 2019 x |
| 2020 1 |
| 2021 x |
| 2022 x |
| 2025 x |
| 2030 x |
| |

| Upper bound of 90% confidence interval |
|--|
| 2010 x |
| 2015 x |
| 2018 x |
| 2019 x |

2030 x

Please provide your estimates for the median and lower and upper bounds of 90% confidence intervals for the cost-effectiveness of the marginal funding opportunity for each of the years listed, by copy-pasting the template above and replacing every 'x' with a numerical value.

[answer]

Please use this space to explain the main reasoning and evidence underlying your estimates

[answer]

Section 3: Follow-up

Please use this space for any other questions, comments or

feedback on this survey

If you would like a direct answer, please reach out to sjir@founderspledge.com

[answer]

Would you be happy to be contacted with any follow-up questions,

based on your answers?

[Yes OR

No OR

Other]

We plan to use this research for a report on investing as a

philanthropic strategy, which will be published on our website.

Would you like to receive the final report at the email address you

specified?

[Yes OR

No OR

Other]

Would you like your name to be mentioned in the

acknowledgements section of the report?

[Yes OR

No OR

Other]

Appendix IV: Analysis of survey results

We shared the survey with a collection of ~20 organizations and individuals that we are aware of with expertise on evaluating the cost-effectiveness of long-term-oriented funding opportunities and/or the availability of such opportunities in the future and our ability to detect them.

Respondents

We received 11 responses of which 6 were external to the Founders Pledge research team. Of the respondents, we classified 4 as subject area experts, based on our existent knowledge of their relevant background and expertise and the answers they provided to the 'background and expertise' question in the survey. The assignment into this expert category was necessarily subjective, and was done by the author with the input of two other Founders Pledge researchers. The author excluded himself from consideration in order to avoid his views being overrepresented in the analysis.

The other respondents also had at least some relevant background, for instance as long-term impact focused funding opportunity researchers, but evaluation of the cost-effectiveness of funding opportunities over time was less at the forefront of their profession. We still included their responses in our analysis because (1) this is a novel and difficult topic with only very few true experts, (2) it is difficult to identify who the true experts are and we may err at this and (3) we suspect that people will hold a large variety of views, and we want to capture this diversity of thinking to get a better sense of what would be an adequate measure of uncertainty around the changes in cost-effectiveness

factor.

See below for an alphabetical list of the respondents, with those we classified as subject area experts underlined:

- Aidan Goth, Founders Pledge
- Carl Shulman, Future of Humanity Institute
- Howie Lempel, 80,000 Hours
- Jaime Sevilla, PhD student at University of Aberdeen
- Johannes Ackva, Founders Pledge
- John Halstead, Founders Pledge
- Max Daniel, Future of Humanity Institute
- Patrick Kaczmarek, Effective Giving
- Sjir Hoeijmakers, Founders Pledge
- Stephen Clare, Founders Pledge
- <u>William MacAskill, Global Priorities Institute</u>

Please note that all of their responses were made on an individual basis, and do not necessarily represent the views of their employers.

Limitations

First of all, we are aware of the extreme gender, geographic and ethnic bias in our sample. This is a consequence of both the overrepresentation of white UK-resident males in this field more generally, and a further overrepresentation in our survey response. For instance, we directly reached out to (only) three women, of whom unfortunately none responded. We do not know if a better representation would have altered the results, or in which direction if so, but we certainly see this as a limitation, and see surveying a more diverse group of experts as an important direction for future research.

Secondly, there are risks of selection and nonresponse bias. For instance, it could be that people who are more optimistic about the prospects of investing to give were more likely to be asked by us and respond than people who are more pessimistic. We had to choose which organisations to survey, and when we shared the survey with these organisations, we had only limited influence over which individuals within those organisations took it. Furthermore, the overall response rate among the people it was ultimately shared with was lower than 50%. This suggests this is a serious limitation, and one that future surveys like these could improve upon. We try to at least partially account for it by not taking simple averages across the estimates, which would overweight potentially overrepresented views. Instead, we analyze the diversity of responses and underlying thoughts, and we take outliers more seriously, to come to an all-considered view ourselves.

Thirdly, our choice to ask for estimates of *relative* rather than *absolute* marginal cost-effectiveness and our way of defining relative marginal cost-effectiveness were - in hindsight - flawed. In our survey instructions, we defined relative marginal cost-effectiveness as follows:

"We ask you to provide your estimates in relative terms to the actual cost-effectiveness of the marginal funding opportunity in 2020. Whatever it may be in reality, here that value is defined and fixed as '1'. Hence, if you think the cost-effectiveness of a marginal funding opportunity in a particular year is 'y' times that of the one in 2020 (for some specified number 'y'), please fill in 'y' for your estimate."

This definition is adequate if one's estimates of the relative marginal cost-effectiveness in past and future years are largely independent of one's view on the marginal cost-effectiveness in 2020 in *absolute* terms, i.e. as expressed as some measure of value added to the world per dollar. However, if those correlate in some way, which they likely do, the definition becomes difficult to interpret, ambiguous or even ill-defined. This is the case, for instance, if one thinks - as some of our respondents do - that the current marginal cost-effectiveness has a meaningful probability of being negative in absolute terms: knowing whether it is currently negative would likely change one's estimate for the future *relative* marginal cost-effectiveness has a meaningful probability of being very close to zero in absolute terms: when it is, a very large future relative marginal cost-effectiveness might seem more likely, as even small changes in absolute cost-effectiveness could lead to large relative changes.¹²⁶

The correlation can play a role both *within* an individual's estimates of a median and lower and upper bounds - if one has a lot of uncertainty about what the current absolute marginal cost-effectiveness is - and *across* the estimates of individuals - if those differ because these individuals have varying views on the current absolute marginal cost-effectiveness.¹²⁷

In addition to and as a consequence of these inherent flaws in our definition, some respondents interpreted our instructions differently from intended and others consciously

decided to provide estimates by a different definition, so as to better be able to convey their views. They all used one of the following three definitions:

Definition 1 (the one originally intended)

- The scale is indexed on the marginal cost-effectiveness in 2020
- '1' represents the marginal cost-effectiveness in 2020, and all other estimates are indexed on this value
- For instance, if the 90% LB in 2030 is 0.07 that means one thinks it is 95% likely that the marginal cost-effectiveness in 2030 is more than 0.07 times the true marginal cost-effectiveness in 2020.

Definition 2

- The scale is indexed on the median estimate of the marginal cost-effectiveness in 2020
- '1' represents the median estimate of the marginal cost-effectiveness in 2020, and all other estimates are indexed on this value
- For instance, if the 90% LB in 2030 was 0.07 that means one thinks it's 95% likely that the marginal cost-effectiveness in 2030 is more than 0.07 times your median estimate of the marginal cost-effectiveness in 2020.

Definition 3

- The scale differs for median, LB and UB estimates.
- For median estimates, '1' represents the median estimate of the marginal cost-effectiveness in 2020. For LB estimates, '1' represents the LB estimate of the

marginal cost-effectiveness in 2020. For UB estimates, '1' represents the UB estimate of the marginal cost-effectiveness in 2020.

 For instance, if the 90% LB in 2030 is 0.07 that means one thinks it's 95% likely that the marginal cost-effectiveness in 2030 is larger than 0.07 times the 90% LB estimate of the marginal cost-effectiveness in 2020.

Fortunately, median estimates mean largely¹²⁸ the same thing across these definitions. Moreover, we were able to infer which definition each of the respondents used,¹²⁹ and combining this with their lower and upper bound estimates and reasoning allowed us to conduct our analysis largely as intended. However, the different definitions do mean that quantitative estimates for lower and upper bounds across responses aren't always comparable. Hence, rather than reporting aggregate statistics for those estimates, we decided to evaluate each response in the context of the definition used and explanations provided, and report the broad takeaways from this analysis for the difference in cost-effectiveness factor.

Results

The graphs below show the median marginal cost-effectiveness estimates relative to 2020 for each of the respondents for the periods of 2010 to 2020 and 2020 to 2030 respectively, each first with a log scale and then with a linear scale on the vertical axis.

We decided to include graphs with both a log and with a linear scale because both have their advantages for interpretation. The log scale fits better the intuition of a multiplicative factor and the idea that '1' is the point of stable marginal cost-effectiveness over time. However, we ultimately care about the *expected* relative marginal cost-effectiveness, and that is better represented on a linear scale. If, for instance, there were a 50% chance of the relative marginal cost-effectiveness in 2030 being 10 and a 50% chance of it being 0.1, this would imply the expected relative marginal cost-effectiveness would be ~5, and hence strengthen the case for investing to give in ten years. A linear scale would visually represent this takeaway whereas on a log scale, '1' is the midpoint between '10' and '0.1'.

For proper resolution and visual comparison across the graphs with the same scale, we decided on the same vertical scale (from 0.1 to 10). This means that a few data points for the 2010 to 2020 graphs are not included. However, this concerns only 8 points (3 for 2015 and 5 for 2010), and these aren't important for observing the overall trends. The absence of a line going left from one of the 2018 data points in the log scale graph is explained by the relevant respondent having estimated their relative marginal cost-effectiveness to be '0' in both 2015 and 2010, which does not show up on a log scale.

Figure IV.1. Graphs with median marginal cost-effectiveness estimates 2010-2030



Median marginal cost-effectiveness estimates 2010-2020, log scale

Experts are in green; other respondents in yellow

Median marginal cost-effectiveness estimates 2020-2030, log scale



Experts are in green; other respondents in yellow

Median marginal cost-effectiveness estimates 2010-2020, linear scale

Experts are in green; other respondents in yellow



Median marginal cost-effectiveness estimates 2020-2030, linear scale

Experts are in green; other respondents in yellow



There are a few key takeaways from these data and the supporting reasoning provided by the respondents:

- Estimates from 2010 and 2015 vary widely, mostly reflecting the different personal histories of the respondents. They converge more in 2018 and 2019.
- Nearly all (10/11) survey respondents have a median estimate of >=1 for 2030, implying that they think that in at least 50% of cases there would be more cost-effective marginal funding opportunities in 2030 than there are now.
- The future estimates of the 4 experts who took the full survey vary widely (more so than those of the other respondents), and this variance becomes more expressed over time: they estimate medians of 0.1, 1.5, 8 and 10 in 2030. However, as evidenced by their explanations, their reasoning does not differ as much as their estimates seem to suggest:
 - The first expert estimates a consistently decreasing marginal cost-effectiveness from 10 in 2010 to 0.1 in 2030. According to their explanation, this is driven mostly by a decreased availability of opportunities via both strongly diminishing marginal returns and increased community spending. They think this will not be outpaced by exogenous learning.
 - The second expert estimates decreasing marginal cost-effectiveness in the past year and coming few years and explains this in the same way. However, they expect exogenous learning to outpace this effect closer to 2030.
 - The third observes the same effect of diminishing returns and increased community spending but thinks exogenous learning and a growing

ecosystem of long-term oriented funding opportunities will already weigh stronger than increased community spending in 2021.

• The fourth again refers to increased spending and diminishing returns as two important factors to consider, but thinks that they will be outpaced by exogenous learning and have been outpaced by that over the past 10 years as well.

Other relevant considerations mentioned by respondents included:

- An increase of socio-technological and political risks over the next few years, increasing the availability of opportunities
- More human resources entering the long-term impact space, increasing the availability of opportunities
- A lag in the effects of those human resources entering the space compared to the effects of new funding becoming available, and less expected extra funding entering the space compared to human resources (relative to their presence now), both increasing the availability of opportunities
- Extrapolating an exponential increase in marginal cost-effectiveness over the past
 10 years which, mostly driven by exogenous learning
- A continuing trend of more targeted marginal funding opportunities
- An increased ability to identify smaller high-impact funding opportunities
- An increased ability to provide complementary resources to funding opportunities when funding them, driving up their cost-effectiveness
- Exogenous learning by nonprofits in a cause area driving an increase in the availability of more targeted, higher-impact opportunities

Note that these were all reasons for the marginal cost-effectiveness to increase over time. What seems clear from the responses is that the main counteracting consideration is increased availability of funding combined with diminishing returns: the differences between respondents' future estimates seem to be mainly explained by how heavily they weigh this consideration against the many potential others mentioned above.

The respondents' 90% lower and upper bound estimates are consistent with this conclusion and offer a few further insights into the respondents' reasoning:

- At least 4 out of 11 respondents, of which 3 experts, estimated a negative lower bound for all years, meaning that they think there is a >=5% probability that the current marginal funding opportunity (and the one in past and future years) has a negative impact on the world. The view of the other 7 on this is unclear, because of the use of different definitions for the marginal cost-effectiveness scale.
- The credible interval bounds become wider from 2020 to 2030 for all respondents but one. The exception is the expert who also estimates a consistent decrease in the median marginal cost-effectiveness, and even though they do not address the narrowing credible interval directly in their explanation, it seems that their decreasing median estimate explains it: if we measure the credible interval width relative to the median estimate for the same year, their 2030 credible interval is wider than their 2020 credible interval.
- Among all 10 respondents other than that expert, there is a clear trend of their lower bound estimates decreasing less than their upper bound estimates are increasing over time, and often to a large extent.

- 5 of these respondents' lower bound estimates in 2030 are roughly the same as their lower bound estimates for 2020, 5 differ by a factor of ten or more, and only 1 of those differs by a factor of a hundred.
- All of their upper bound estimates in 2030 differ by a factor of three or more from their upper bound estimates for 2020, 7 differ by a factor of ten or more, 3 differ by a factor of a hundred or more, and two by a factor of a thousand.
- Looking at these differences in absolute terms paints an even clearer picture that the distribution they have in mind has a positive expected difference in the cost-effectiveness factor over ten years, and in most cases a much larger one than considering only their median estimates would have suggested.

However, we cannot derive a credible interval for the difference in cost-effectiveness factor directly from the respondents' 90% credible interval estimates. This is in part because of the different definitions respondents used. But even if they had all used definition 1, it would still not have been appropriate to take their estimates literally.

The issue is that one's estimates for the *relative* marginal cost-effectiveness in 2030 are arguably negatively correlated with the current marginal cost-effectiveness in *absolute* terms. For instance, *given* that the current marginal cost-effectiveness is low or even negative in absolute terms, it seems more likely for marginal cost-effectiveness to substantially increase in the next ten years, so that one should estimate a larger *relative* marginal cost-effectiveness in 2030, because there is a lot more room for improvement. And the other way around: *given* current marginal cost-effectiveness already being very high in absolute terms, one should probably estimate a smaller *relative* marginal

cost-effectiveness in 2030, for example because there will be fewer crucial insights left to obtain via exogenous learning.

This means that one's upper bound estimate for the relative marginal cost-effectiveness in 2030 is more likely to reflect the cases in which current marginal cost-effectiveness is assumed to be high in absolute terms, whereas one's lower bound estimate is more likely to reflect cases in which current marginal cost-effectiveness is assumed to be low. And that, in turn, changes the calculation of our *expected* marginal cost-effectiveness in ten years from now, in both absolute and relative terms. A perfect negative correlation could go as far as cancelling the effect that the chance of an 'x'-fold increase usually weighs more heavily in expected value terms than an equal chance of an 'x'-fold loss.

To clarify this, we'll consider a toy example. Say we think it is 50/50 whether the marginal cost-effectiveness in 2020 is '1' or '10' in some *absolute* unit of cost-effectiveness, e.g. some fixed amount of reduction in existential risk per dollar. Now say that we think that it is also 50/50 whether marginal cost-effectiveness will increase by a factor of 10 or decrease by a factor of 10 over the next ten years. That is, our 2030 90% credible interval for relative marginal cost-effectiveness is [0.1,10]. The table below shows how a negative correlation as specified above changes the *expected* relative marginal cost-effectiveness in ten years, which is what we are ultimately after.

| Situation | 2020 expected marginal cost-effectiveness (unit of existential risk reduction / dollar) | 2030 expected marginal cost-effectiveness (unit of existential risk reduction / dollar) | 2030 expected relative marginal cost-effectiveness (indexed on 2020 expected marginal cost-effectiveness) |
|---|--|--|--|
| No correlation between 2020 marginal cost-effectiveness and 2030 relative cost-effectiveness | 5.5 (0.5 * 1 + 0.5 * 10) | 27.775 (0.25 * 0.1 + 0.25 * 1 + 0.25 * 10 + 0.25 * 100) | 5.05 (27.775 / 5.5) |
| Perfect negative correlation between 2020 marginal cost-effectiveness and 2030 relative cost-effectiveness | 5.5 (0.5 * 1 + 0.5 * 10) | 5.5 (0.5 * 10 + 0.5 * 1) | 1 (5.5 / 5.5) |

For the same [0.1,10] 90% credible interval, the case with no correlation yields a more than five-fold larger expected relative marginal cost-effectiveness here than the case with perfect negative correlation. This illustrates how, if we were to just calculate our expected relative marginal cost-effectiveness directly from fitting the respondents' reported credible intervals to probability distributions - without accounting for the possibility of negative correlations - we would end up with an expected relative marginal cost-effectiveness that would be heavily biased upwards.

Overall estimate

So, if not by looking at the estimates of the 90% credible intervals directly, how do we estimate a suitable credible interval based on the survey responses? The approach we have taken is to consider the variation in the predicted ten-year *median* estimates *across respondents* as a starting point, as those are less sensitive to correlations with the current marginal cost-effectiveness.¹³⁰ We then adjust based on the qualitative takeaways from our respondents' 90% credible interval estimates.

Concretely, we start with the most extreme median estimates on either side, 0.1 and 10, as the upper and lower bounds of our 90% credible interval.¹³¹ From here, we should arguably increase the upper bound, because it is clear from the respondents' bound estimates that they judged the probability of large increases (>=100-fold) in marginal cost-effectiveness as much higher than that of large decreases. On the other hand, we want to avoid overweighting large increases by failing to take into account a negative correlation with respondents' views on current cost-effectiveness *across* respondents estimates, and we want to give some extra weight to the perspective of our expert who thinks marginal cost-effectiveness will decrease over time. Weighing up these considerations, we choose to only double the upper bound to 20.

For our overall median estimate of the difference in cost-effectiveness factor, naively taking an average of the respondents' median estimates for the 2030 relative marginal cost-effectiveness would yield an estimate of 3.8 when considering all respondents and 4.9 when considering experts. However, that would overweight large outliers. Furthermore, it could be biased upwards as it does not account for negative correlations between each
respondents' view of current marginal cost-effectiveness in 2020 in absolute terms and their 2030 median relative marginal cost-effectiveness estimate. And finally, as mentioned above, such an estimate would be highly sensitive to selection and nonresponse bias.

Instead, we want to give proper weight to the one expert who predicts a median tenfold decrease in marginal cost-effectiveness, to the fact that most respondents predict only only a moderate median increase, and to the hypothesis that the respondents who predicted larger increases think current marginal cost-effectiveness is lower. We hence go with something close to a median of the respondents' median estimates, adjusted downwards for the possibility of selection and nonresponse bias: 1.2.

Fitting these estimates to a log-normal distribution using a least-squares approximation¹³² yields a final credible interval estimate of approximately [0.09,17] with a median of 1.2 and an average of ~4.5 for the ten-year difference in cost-effectiveness factor.

Endnotes

1. Michael D. Klausner, 'When Time Isn't Money: Foundation Payouts and the Time Value of Money', *SSRN Electronic Journal*, 2003, https://doi.org/10.2139/ssrn.445982.

See <u>this article</u> by career advice organisation 80,000 Hours for a good introduction to general considerations when trying to benefit the long term. We aim to publish more on this ourselves soon.
 Credible intervals are the <u>Bayesian</u> equivalent of confidence intervals. See <u>this page</u> for more on them.

4. This is the sum of all the values the impact ratio could take, weighted by their estimated probabilities.

5. William MacAskill, 'When Should an Effective Altruist Donate?' (Global Priorities Institute, 24 September 2019), 10,

https://globalprioritiesinstitute.org/william-macaskill-when-should-an-effective-altruist-donate/. 6. Shane Frederick, George Loewenstein, and Ted O'Donoghue, 'Time Discounting and Time Preference: A Critical Review'. *Journal of Economic Literature* 40. no. 2 (2002): 351-401.

7. Moritz A. Drupp et al., 'Discounting Disentangled', *American Economic Journal: Economic Policy* 10, no. 4 (November 2018): 109–34, https://doi.org/10.1257/pol.20160240.

8. Charles-Joseph Mathon de La Cour, *Testament de M. Fortuné Ricard, maitre d'arithmétique à D***, [par Mathon de la Cour] lu & publié à l'audience du bailliage de cette ville, le 19 août 1784. [Suivi des Tables justificatives],* 1784.

9. Klausner, 'When Time Isn't Money'.

10. From inputting 1791, £1000, and 2019 in <u>this calculator</u>. Note that inflation adjustments are somewhat controversial and should be taken with a grain of salt, certainly over long timescales like these. For more on that, see e.g. <u>this working paper</u>.

11. Klausner, 'When Time Isn't Money'.

12. We cannot say with certainty that Franklin has had more impact by investing to give than he would have had if he had spent the money at the time, but his case at least demonstrates success at the financial side (even in inflation-adjusted terms) of investing to give over long timescales.

13. There may have also been some benefits to fixing these things, in that they may have helped increase the persistence of the funds (see Section 2.3 for a discussion of this factor). For instance, fixing the beneficiaries to be Philadelphia and Boston might have made it less likely for Franklin's will to be challenged in court and easier to allocate the funds when they were released. But these benefits are unlikely to outweigh the substantial costs.

14. We will primarily examine investing to give from the perspective of maximising our impact over the long term, as this would be the objective of the Fund. In Section 2.6 and 2.7, we will briefly consider its promise from the perspective of those who aim to benefit the current generation and those who aim to avert animal suffering.

15. This is a crucial difference with considering to invest or spend now for personal gain: in the latter case you are the only one who is immediately spending on yourself or investing, and so it does make sense to diversify your strategy.

16. This is arguably true of <u>Open Philanthropy</u>, who spend <u>on the order of \$100 million per year</u> on funding opportunities to maximise impact from a long-term perspective, and plan to spend billions more over the course of the following decades.

17. For instance, at Founders Pledge we publish nearly all of <u>our research</u> online, both so it can be independently critiqued and so it can be used by others.

18. However, if we choose to set up such a Fund, we will certainly aim for it to do so, and we plan to re-examine the diversification question when it becomes more relevant.

19. Philip Trammell takes this approach in <u>this paper</u> (in progress): he provides a model that approximates the optimal community portfolio under certain assumptions. We think this is an approach that is worth pursuing and might lead to practicable insights in the future. However, for

our present practical purposes we think the assumptions necessary in the model, e.g. defining a particular diminishing marginal returns function over all time, are too limiting.

20. We chose ten years for our proxy question in particular to balance (1) avoiding considerations that are very local in time and not relevant to this question more generally, e.g. the current coronavirus pandemic, (2) avoiding the increased difficulty that comes with trying to make concrete predictions over longer timelines, (3) being able to test our predictions in the coming few years already, and (4) availability of data when looking back over this time horizon now.

21. See Section 3.1.1 for a discussion of investment-like giving opportunities.

22. We discuss this factor, which we call "exogenous learning", in Section 2.4.2.

23. For more on index funds, see this page.

24. See <u>this article by 80,000 Hours</u> for a good introduction to general considerations when trying to benefit the long term; we aim to publish more on this ourselves soon.

25. See Section 3.1.1 for a definition and discussion of investment-like giving opportunities.

26. Our work here substantially borrows from and builds on work by other researchers, for instance William MacAskill's <u>qualitative exploration of the question when to give</u> and Phil Trammell's <u>paper on</u> <u>patient philanthropy</u>.

27. We are implicitly assuming independence among the factors here. We discuss the implications of this limitation in Section 3.3.1.

28. These are the <u>Bayesian</u> equivalent of confidence intervals. See <u>this page</u> for more on them.

29. This is motivated by application of the <u>central limit theorem</u> in the log domain. See <u>this page</u> for more on the properties of log-normal distributions.

30. See <u>this page</u> for more on the properties of heavy-tailed distributions.

31. See <u>this page</u> for more on the properties of beta distributions.

32. For more on Monte Carlo simulation, see <u>this page</u>.

33. For more on value drift, see <u>this page</u>.

34. "For example, the least effective HIV/AIDS intervention produces less than 0.1 percent of the value of the most effective. In practical terms, this can mean hundreds, thousands, or millions of additional deaths due to a failure to prioritize."

'The Moral Imperative toward Cost-Effectiveness in Global Health', Center For Global Development, accessed 9 October 2020,

https://www.cgdev.org/publication/moral-imperative-toward-cost-effectiveness-global-health.

"It's plausible that the differences in cost-effectiveness between social programs can vary 35. For more on this, please see Section 3.1.1 on investment-like giving opportunities.

36. An important consideration in the design of a Fund is the (seeming) trade-off between preventing value drift and allowing for potential value improvement over time. An example of a way to balance those objectives would be to appoint value-aligned trustees that select their own successors.

37. See <u>this</u> and <u>this</u> paper by William MacAskill for a qualitative overview of the most relevant considerations.

38. Please see Appendix IV for an explanation of how we arrived at these particular estimates. 39. As explained in our section on the model, Guesstimate has a sampling limit of 5000 per variable, causing the model to give slightly different results on each run. However, running the model 20 times never yielded a median that was smaller or larger than 2.0 or 2.3 respectively, and never yielded an average that was smaller or larger than 8.2 or 11 respectively, so the sampling limit doesn't meaningfully affect our results.

40. Technically, she could spend the money on opportunities which turn out to have a negative impact, but she would never choose to spend it on opportunities with a negative *expected* impact. This means that the only scenario in which investing to give would have a negative impact is the one in which we would be reliably wrong about the sign (positive/negative) of our future expected impact estimates. However, even if this were the case, it seems likely that our current estimates would be even more inaccurate. On the other hand, we should note that the stakes in the future could be higher, for example through us having access to more powerful technology. 41. For example the fact that interest rates are currently exceptionally low

42. We say largely because the persistence of funds is likely somewhat correlated with one's aims. It could for instance be that it is easier to value drift from certain values than others, or that funds aimed at certain goals face more external legal challenges than others. However, we expect these effects to be negligible - or infeasible to forecast with current data insofar as they aren't negligible - at least for the timescale of ten years.

43. <u>GiveWell</u> is a nonprofit dedicated to finding outstanding giving opportunities and publishing the full details of their analysis to help donors decide where to give. Unlike charity evaluators that focus solely on financials, assessing administrative or fundraising costs, they conduct in-depth research aiming to determine how much good a given program accomplishes (in terms of lives saved, lives improved, etc.) per dollar spent. Rather than try to rate as many charities as possible, they focus on the few charities that stand out most (by <u>their criteria</u>) in order to find and confidently recommend high-impact giving opportunities (<u>their list of top charities</u>).

44. 'Top-Rated Charities - 2010 Archived Version', GiveWell, accessed 9 October 2020, https://www.givewell.org/charities/top-charities/2010.

45. 'Our Top Charities | GiveWell', accessed 9 October 2020,

https://www.givewell.org/charities/top-charities.

46. 'Top Charities - November 2011 Archived Version', GiveWell, accessed 9 October 2020, https://www.givewell.org/charities/top-charities/November-2011.

47. 'Our Top Charities | GiveWell'.

48. 'Update on No Lean Season's Top Charity Status', The GiveWell Blog, 19 November 2018, https://blog.givewell.org/2018/11/19/update-on-no-lean-seasons-top-charity-status/.
49. This is excluding potential compound social returns from giving earlier: those are discussed in Section 3.1.1.

50. We should note that GiveWell <u>doesn't just rely</u> on these cost-effectiveness estimates for making their recommendations, but comparing them across time still serves as a good proxy for how cost-effectiveness has changed all-things-considered, especially given that we'd expect little change in the other factors GiveWell takes into account (e.g. an organization's track record) if we are comparing the same organization over time.

51. 'Approaches to Moral Weights: How GiveWell Compares to Other Actors', GiveWell, accessed 9 October 2020,

https://www.givewell.org/how-we-work/our-criteria/cost-effectiveness/comparing-moral-weight s.

52. As a sense-check, we also had a quick look at how the

cost-per-death-of-an-under-5-year-old-averted-equivalent of these charities evolved: for AMF this went from \$2000 in 2012 to \$1700 by the end of 2019, and for SCI from \$1500-\$4000 in 2012 to \$2900 in 2020, suggesting a ten-year difference in cost-effectiveness factor of close to '1'.

https://blog.givewell.org/2012/12/19/cost-effectiveness-of-nets-vs-deworming-vs-cash-transfer s/

https://docs.google.com/spreadsheets/d/1zLmPuddUmKsy3v55AfG_e1Quk-ngDdNzW-FDx0T-Y94 /edit#gid=1034883018

53. Holden, 'Cost-Effectiveness of Nets vs. Deworming vs. Cash Transfers', The GiveWell Blog, 19 December 2012,

https://blog.givewell.org/2012/12/19/cost-effectiveness-of-nets-vs-deworming-vs-cash-transfer s/.

54. '2019 GiveWell Cost-Effectiveness Analysis – Version 6 (Public) - Google Sheets', accessed 9 October 2020,

https://docs.google.com/spreadsheets/d/1zLmPuddUmKsy3v55AfG_e1Quk-ngDdNzW-FDx0T-Y94 /edit#gid=1034883018.

Both estimates are in nominal dollars at the time of the analysis.

55. Again, these figures are in nominal terms.

We used the returns between Jan 1 2010 and 31 December 2019 from this online tool.

Also note that as of 22 September 2020, in spite of COVID, the S&P 500 has gone <u>up by 1.6%</u> over 2020 so far.

56. This is an important caveat: our example does not consider the expected impact *in 2010* of investing to give in ten years vs giving in 2010 here, but the expected impact *in 2020*. However, it seems likely that we would have found the expected impact in 2010 of investing to give in ten years to be similar or higher: even though expected financial returns would have been lower, expected exogenous learning would have been higher, given the short existence of high-impact evaluators like GiveWell and how much they had already changed their views and recommendations in the first few years. The fact that AMF is still among the highest-impact funding opportunities today is arguably somewhat surprising from a 2010 perspective.

57. We are again assuming limited funds available for donation. If a donor e.g. instead had \$10 billion available to give to AMF in 2010, this would have led to diminishing marginal returns and would have likely meaningfully changed the trajectory of the organization, so that the cost-effectiveness estimates in both 2010 and 2020 would no longer have been valid.

58. Email correspondence with GiveWell, 6 October 2020. This research was conducted by GiveWell Senior Fellow Alex Cohen. The full research publication is forthcoming on <u>GiveWell's website</u>.

59. GiveWell accounted for ~\$140 million moved to their recommendations in <u>their 2018 metrics</u> report, and growing the amount of money moved each year is a priority of theirs.

60. See the section on inflation in Appendix I for our reasons to use a 2% projected inflation rate. 61. A caveat here is that both GiveWell and Founders Pledge have only recently started looking more extensively into policy advocacy interventions. We think it more likely than not that these explorations will yield new and potentially higher-impact funding opportunities in the coming few years.

62. For instance, if GiveWell or Founders Pledge discovers a new funding opportunity that is more than twice as cost-effective as our current recommendations.

63. 'Farmed Animal Funders', accessed 9 October 2020, https://farmedanimalfunders.org/.

64. 'Animal Charity Evaluators | Helping People Help Animals', accessed 9 October 2020, https://animalcharityevaluators.org/.

65. Email correspondence with Farmed Animal Funders, 28 September 2020.

66. We discuss compound returns on giving and investment-like giving opportunities in Section 3.1.1.

67. Email correspondence with Animal Charity Evaluators, 1 October 2020

68. For instance in <u>this blog post</u> by GiveWell.

69. The technical reasons for this are discussed in more depth in Section 5.1.2 and 5.1.3 of Philip Trammell's 2020 working paper <u>Discounting for Patient Philanthropists</u> and Paul Christiano's 2013 blog post <u>Giving now vs. later</u>. Robin Hanson makes a separate theoretical argument against large compound returns on giving in <u>this blog post</u>.

70. World Bank, 'GDP Growth (Annual %) | Data', accessed 9 October 2020,

https://data.worldbank.org/indicator/NY.GDP.MKTP.KD.ZG?end=2019&start=2009.

71. World Bank, 'GDP Growth (Annual %) | Data'.

72. For more on these, see <u>our research report on existential risks</u>.

73. In our survey for the difference in cost-effectiveness factor in our model, we additionally restricted this later point to be at least ten years from now. So our model and proxy question technically include investment-like giving opportunities that could not benefit from the full ten years of exogenous learning and hence differ substantially from "investing to give in ten years". 74. For more on the effective altruism movement, see this website.

75. For more on global priorities research, see <u>this article</u> by career advice organisation 80,000 Hours.

76. This is unless it would then be funded by someone who would otherwise have invested to give or given to similarly impactful investment-like giving opportunities.

77. See <u>this page</u> for more on Donor-Advised Funds.

78. As we noted in Section 2.6.2, GiveWell alone accounted for ~\$140 million moved to their recommendations in <u>their 2018 metrics report</u>. In addition, there are lots of other actors spending on similar interventions, for example governments:

"In 2019, official development assistance (ODA) by member countries of the Development

Assistance Committee (DAC) totalled USD 152.8 billion, representing 0.30% of their combined GNI" OECD, 'Aid by DAC Members Increases in 2019 with More Aid to the Poorest Countries', 16 April 2020, https://www.oecd.org/dac/financing-sustainable-development/development-finance-data/ODA-2 019-detailed-summary.pdf.

79. 'Our Progress in 2019 and Plans for 2020 | Open Philanthropy', accessed 9 October 2020, https://www.openphilanthropy.org/blog/our-progress-2019-and-plans-2020.

80. From selecting Farm Animal Welfare as Focus Area in:

81. This is not necessarily true for investment-like giving opportunities, which may allow for reallocation of resources at later points in time.

82. MacAskill, 'When Should an Effective Altruist Donate?', 10.

83. This can be conceptualized as exogenous learning on the level of evaluating investing to give vs giving now, rather than on the level of evaluating giving opportunities. Note that some of the learning can be done endogenously (this research project is an example), but that we should expect exogenous learning to occur as well. For instance, as we get better at evaluating the long-term cost-effectiveness of funding opportunities via the research of others, we should expect to get better at estimating the change in marginal cost-effectiveness over time as well.

84. See this page for more on leveraged investing.

85. See also our discussion of correlations between the financial returns and difference in cost-effectiveness factor in Section 3.3.1.

For more in-depth discussions of the significance of investment risk for altruistic investors, see <u>this</u> <u>article</u> by Paul Christiano and <u>this article</u> by Michael Dickens.

86. John Halstead, 'Impact Investing', Founders Pledge, accessed 9 October 2020,

https://founderspledge.com/stories/impact-investing-executive-summary.

87. Frederick, Loewenstein, and O'Donoghue, 'Time Discounting and Time Preference'.

88. Moritz A. Drupp et al., 'Discounting Disentangled', *American Economic Journal: Economic Policy* 10, no. 4 (November 2018): 109–34, https://doi.org/10.1257/pol.20160240.

89. For a more in-depth treatment of the time preference argument for investing to give, please see Philip Trammell's working paper <u>Patience and Philanthropy</u>.

90. There is an argument to be made that positive correlations of a certain strength have a larger upwards influence on the expected impact than negative correlations of the same strength have a downwards influence, and that a symmetric increase in uncertainty should hence make us slightly more optimistic about investing to give. When we say the correlations increase our overall uncertainty, we have already incorporated this effect into our judgement: we don't think it is strong enough here to tilt the overall balance in favor of investing to give, mostly because we think a strong negative correlation is slightly more likely than a strong positive correlations between the financial returns and difference in cost-effectiveness factors.

91. This is again assuming the Long-Term Investment Fund hasn't grown to an extremely large size (yet). At that point, considerations such as risks of loss of ownership due to legal challenges might become more important and change our expectations here.

92. Interestingly, if one thinks it likely that there are very strong diminishing marginal returns to spending at the community level, this both strengthens one's predicted effect of this negative correlation and makes a decrease in marginal cost-effectiveness generally look more likely, as discussed in section 2.4.3. On the other hand, if one thinks diminishing marginal returns aren't very strong, this weakens the negative correlation and increases one's estimate of marginal cost-effectiveness over time. Hence, one's views on diminishing marginal returns at the community level will likely be quite strongly correlated with how one views the potential of investing to give.
93. This dynamic is also what enables 'mission hedging': investing in companies that could make a particular problem worse if they are highly successful, e.g. oil companies, so that the investor makes a profit whenever the problem is at its worst and there is the highest need for funding solutions.
94. For our views on how to best tackle climate change as a philanthropist, see <u>our report on climate change</u>.

95. As we noted in Section 2.6.2, GiveWell alone accounted for ~\$140 million moved to their recommendations in <u>their 2018 metrics report</u>, and as we noted in Section 3.1.2, global government aid spending is more than 100 billion per year:

96. From eyeballing their grants database.

97. See Section 2.7.1.

98. Philip Trammell, 'Patience and Philanthropy', n.d., 39.

99. For more on what annualized returns are, see this page.

100. Calculating annualized real returns using <u>this online tool</u> yields 7.08% for the period 1 January 1927 to 31 December 2019.

101. 'Rio de Janeiro Stock Exchange', in Wikipedia, 31 July 2020,

https://en.wikipedia.org/w/index.php?title=Rio_de_Janeiro_Stock_Exchange&oldid=970522311. 102. "Over the 120 years from 1900 to 2019 [...] the real return on the world index was 5.2% per year for equities [...] Most recently, we have added Austria, Portugal, China and Russia. Austria and Portugal have a continuous history, but China and Russia do not. To avoid survivorship bias, all these countries are fully included in the world indexes from 1900 onward."

103. This is from trying out several distributions in Guesstimate by hand and eyeballing the data. The final match can be found <u>here</u> below our model for the ten-year impact ratio. Note that, as we annualize only over ten-year periods here, the 1.9 mean factor of our log-normal distribution is higher than the 1.7 factor that would result from a constant 5.2% yearly rate of investment returns.
104. Dimson, Marsh, and Staunton, *Credit Suisse Global Investment Returns Yearbook 2020*. Table 1, p. 23

105. Calculating the standard deviation of real returns using <u>this online tool</u> yields 20.07% for the period 1 January 1927 to 31 December 2019.

106. Bryan Caplan, 'Paasche Says Progress', Econlib, 28 March 2019,

https://www.econlib.org/paasche-prosperity/.

107. See Section 2.4 and Appendix IV

108. We did not choose to leave it in in the first place, because the data we had available that was adjusted for survivorship and selection bias was inflation-adjusted as well.

109. David Wessel, 'Alternatives to the Fed's 2 Percent Inflation Target', *Brookings* (blog), 7 June 2018, https://www.brookings.edu/research/alternatives-to-the-feds-2-percent-inflation-target/.
110. World Bank, 'Inflation, Consumer Prices (Annual %) - United States | Data', accessed 9 October 2020, https://data.worldbank.org/indicator/FP.CPI.TOTL.ZG?locations=US&name_desc=false.
111. World Bank, 'Inflation, Consumer Prices (Annual %) | Data', accessed 9 October 2020,

https://data.worldbank.org/indicator/FP.CPI.TOTL.ZG?name_desc=false.

112. For more on the effective altruism movement, see this website.

113. Benjamin Todd, 'More Empirical Data on "value Drift" - EA Forum', accessed 9 October 2020, https://forum.effectivealtruism.org/posts/eRQe4kkkH2pPzqvam/more-empirical-data-on-value-dr ift.

114. Todd.

115. Indeed, the current research project's main aim is to explore whether - and if so, how - we should set up such a Fund to support Founders Pledge members in maximizing their impact.

116. Todd, 'More Empirical Data on "value Drift" - EA Forum'.

117. For more on global catastrophic risks, see this page.

118. For more on existential risks, see <u>our report</u> on the topic.

119. Toby Ord, *The Precipice: Existential Risk and the Future of Humanity*, Illustrated edition (New York: Hachette Books, 2020).

120. Anders Sandberg, *Grand Futures: Visions and Limits of What Can Be Achieved*, Forthcoming. 121. For examples, see: 'Trust Issues | Paul Collins', Lapham's Quarterly, accessed 6 August 2020, https://www.laphamsquarterly.org/future/trust-issues.

122. 'Sector Overview', accessed 9 October 2020,

https://register-of-charities.charitycommission.gov.uk/sector-data/sector-overview.

123. 'Top 10 Charities', accessed 9 October 2020,

https://register-of-charities.charitycommission.gov.uk/sector-data/top-10-charities.

124. 3%: Teresa Harrison and Christopher Laincz, 'Entry and Exit in the Nonprofit Sector', *Contributions to Economic Analysis & Policy* 8 (1 February 2008): 1937–1937, https://doi.org/10.2202/1935-1682.1937.

5%: Jesse Lecy and Eric Joseph van Holm, 'The End Is Nigh: Limits to Growth of the Nonprofit Sector', *SSRN Electronic Journal*, 1 January 2013, https://doi.org/10.2139/ssrn.2478238. 125. We did this by trial and error in <u>our Guesstimate model</u>.

126. One seeming way to deal with this is to express current marginal cost-effectiveness in terms of a probability distribution rather than a point estimate, but dividing by this probability distribution is also ill-defined or at least ambiguous. Alternatively, one could index on the expected value of this probability distribution, but that would no longer capture any information on the correlations and hence distort our final estimates.

127. We discuss how we adjust our estimates to account for these in the section on results below. 128. Technically, the median estimates under definition 1 are different from those under definition 2 and 3. However, based on the respondents' explanations we don't expect this difference to have meaningfully affected their estimates, and certainly not to have affected the general trends in their estimates.

129. We inferred this indirectly from their estimates and explanations and directly by following up with them via email.

130. Doing it this way avoids the effects of such correlations where they exist within any individual's estimates, but they can still have an effect across people's views, as people will have different (median) views of the current marginal cost-effectiveness and will base their ten-year marginal cost-effectiveness estimate on those. This is another reason to use the median estimates rather than the credible interval estimates people provided: we are not interested in the spread of the difference in cost-effectiveness conditional on one particular view of current marginal cost-effectiveness, but in the spread of the difference in cost-effectiveness accounting for all the values current marginal cost-effectiveness could take, which, because of the negative correlation, is likely less than that of the former.

131. Conveniently, these are the same whether we consider only our labeled experts or all respondents.

132. We used <u>this online tool</u> to do this