

CHARITY ENTREPRENEURSHIP

How to Start a **HIGH-IMPACT NONPROFIT**



CONTRIBUTORS INCLUDE:

JOEY SAVOIE, PATRICK STADLER, ISHAAN GUPTASARMA

Acknowledgments

This handbook couldn't have been achieved without the help and support of numerous contributors. We'd like to take a moment to thank some of the individuals who so kindly offered their time.

We are profoundly grateful to all our guest writers for lending their expertise. Our thanks to Peter Brietbart, for his thoughts on going for-profit; to Aaron Hamlin, for his wonderfully comprehensive discussion of starting a nonprofit from scratch; and to Varsha Venogupal, for her advice on corporate fundraising. Thanks also to Kat Woods and Peter Hurford, whose contributions to CE and our umbrella organization Charity Science Foundation have helped so much over the years.

Within CE, many thanks go to Erik Hausen, Juliette Finetti, Andrew Player, Antonia Shann, and Urszula Zarosa. Their knowledge of diverse areas, including monitoring and evaluation, fundraising, operations, and digital media, was invaluable in putting together this handbook, and we're incredibly fortunate to have them on the CE team.

Warmest thanks as well to our wonderful interns for their vital contributions. Thank you to our volunteer copyeditors Lavanya Mane and Rebecca Kirby, for their eagle eyes and their patience in helping out with formatting. Thank you to our graphic designer Nicoleta Faina, for her attention to detail and for creating such a fantastic cover.

Thank you all.

Ishaan Guptasarma, Joey Savoie, & Patrick Stadler.

Table of contents

Introduction: Saving a Million Lives

Part 1: Introduction to Charity Entrepreneurship

I. Advantages

1. The Case for Charity Entrepreneurship
2. Impacts on Your Future Ability to Do Good
3. How Your Team Benefits From Charity Entrepreneurship
4. The Impact of Charity Entrepreneurship on the NGO Sector
5. The Benefits on Top of Changing the World: Four Reasons You Should Consider Charity Entrepreneurship (Besides Impact)

II. Disadvantages

1. Three Years Later
2. What Would Have Happened Otherwise?
3. Six Ways Your Charity Start-up Might Fail, and How to Prevent That
4. Five Challenges a Charity Entrepreneur Faces – and How to Manage Them

III. Assessing fit & preparing

1. What Traits Make a Great Charity Entrepreneur?
2. Which Jobs Will Best Prepare You To Become A Charity Entrepreneur?
3. Should I Start a Charity Now or Later?
4. Plan B: Solid Alternatives to Charity Entrepreneurship
5. How to Increase Your Odds of Getting Accepted to the Charity Entrepreneurship Incubation Program

IV. Co-founder

1. How to Successfully Pick a Co-founder
2. Working with Your Co-founder: How to Excel at Joint Decision-making, Task Management, and Communications
3. How to Strengthen Your Co-founder Relationship

V. Community & Cause Areas

1. Introduction to Effective Altruism
2. The EA Community and Animals
3. The EA Community and Human Interventions
4. The EA Community and Other Interventions
5. Strengths and Weaknesses of Different Cause Areas and Charity Ideas
6. Six Examples of New Charities in Effective Areas

7. I Have a Charity Idea. Any Advice?

Part 2: Making Good Decisions

I. Meta tools (the toolbox)

1. When to Focus and When to Reevaluate
2. The Importance of Time Capping
3. Iterative Depth
4. Useful Databases for Decision-making
5. Charities Love Feedback (As Long As It's Complimentary)

II. Multi-factor decision-making (the Swiss Army knife)

1. Using a Spreadsheet to Make Decisions
2. Why Location Matters When Picking an Intervention
3. Geographic Selection Templates
4. Assessments Made Easy With Google Spreadsheets
5. Weighted Factor Models

III. Tools of rationality (the hammer)

1. What Does It Take to Be A Skeptical Charity?
2. What If You're Working on the Wrong Cause?
3. How CE Uses Informed Consideration

IV. Tools of science (the screwdriver)

1. The Importance of Evidence

V. Tools of effective altruism (the wrench)

1. Counterfactual Impact: What Would Happen If You Didn't Act?
2. The Option Value Paradox
3. Equal Application of Rigor for EA Interventions
4. Why We Look at the Limiting Factor Instead of the Problem Scale

VI. Long-term planning (the pencil)

1. Theory of Change
2. Long-term Planning
3. Going For-profit
4. To Scale, or Not to Scale: That Is the Question?

VII. Problem-solving (the duct tape)

1. Problem-solving in Five Steps
2. How to Come to Better Decisions by Playing Steelman Solitaire

VIII. Measurement & evaluation (the tape measure)

1. Six Examples of Measuring Incomplete Metrics... And How to Fix Them
2. What Is an Impact Evaluation?

3. Monitoring and Evaluation
4. Experimental and Quasi-experimental Design
5. Implementing Surveys in Developing Countries

IX. Task planning (the clamps)

1. Achieving Pareto Productivity: Simple Task Management and Productivity Rules that Go a Long Way

X. Independent experts (the cell phone)

1. How CE Uses Expert Views
2. How Do You Know Who to Trust?

XI. Budgeting & financial planning (the saw)

1. Five Ways to Cut Costs
2. What Is a Reasonable Financial Safety Net for a Charity Entrepreneur?

XII. Cost-effectiveness analysis (the magnifying glass)

1. Introduction to Cost-effectiveness Analysis
2. Quantifying Complex Values
3. 13 Common CEA Mistakes
4. Using Speculative CEAs to Make Decisions
5. Using More Rigorous CEAs: New Incentives, Fortify Health, and CSH
6. How CE Uses Cost-effectiveness Analysis

Conclusion & Summary

Introduction: Saving a Million Lives

Few people have had the opportunity to save someone's life. Those who have, say it's a profound experience to think about a person who would not be alive if not for your actions. And the impact does not stop at the person you save – consider the parents, friends, and many others whose lives are forever enriched by their deep connection to the person you saved. Some people get a single opportunity in their life to save another person, often at great personal risk to themselves, such as running into a burning building. Others take jobs that give them more opportunities to save lives, such as becoming a paramedic. However, few jobs give the opportunity to affect someone's life so drastically.

What about saving 1 million lives? This is a tier reserved for only the most talented scientists, the most well-known political activists, and the most successful charity entrepreneurs. Every great organization which has made a positive difference in the world had to be started by someone. Just as many of the richest people on the planet are entrepreneurs in the for-profit world, many of the most impactful people are entrepreneurs in the nonprofit world.

Saving a million lives is hard for the mind to grasp. Even if you met one person every hour of every day for the rest of your life without sleep, you would not meet a million people. Many cities have a population of a million, and saving a city is a task normally reserved for superhero movies. Yet this feat can be achieved, with luck and a lot of hard work.

The main reason to consider charity entrepreneurship as a career is the possible impact you can have. Historically, people who have founded charities have made a massive difference to countless lives and tackled incredibly important problems that otherwise would have been neglected. The best charities differ from average ones by several orders of magnitude, but if you have an intervention backed by a strong research process and a good set of decision-making tools, becoming a high-impact charity can be an achievable goal. Starting a high-impact organization is often the greatest good that a person can achieve with her/his career. With the right founder, an organization executing a high-impact intervention can make a massive difference to a very large number of lives.

This handbook is about how to become that founder. It outlines how to start that impactful organization, and ultimately how to have the largest positive impact possible on the world.

Part 1: Introduction to Charity Entrepreneurship

I. Advantages

1. The Case for Charity Entrepreneurship

Author:


Ishaan Guptasarma,
Curriculum Developer and Content Writer

*Note: “Charity entrepreneurship” is what you’re about to embark upon as you build your own charity start-up. “Charity Entrepreneurship” (CE) is us – a project of the **Charity Science Foundation** that helps entrepreneurs like you found new effective charities.*

If you’re like us at CE, you want to have the greatest positive impact possible on the causes that prevent suffering and do the most good. For many reading this book you have a wide variety of skills that could be applied in a number of areas, which means that you must compare many possible career options to discover which path will deliver the greatest good to the world.


GiveWell (a well-regarded charity evaluator) suggests using a **cluster approach** when making comparisons between options that are difficult or seem impossible to compare directly. Rather than taking a single sequence of premises and reaching a conclusion (like a math equation), a cluster approach favors conclusions that are robust from

Sequence Thinking



Cluster Thinking

(Larger circle = more robust perspective)



a variety of perspectives and methods of analysis (like a meta analysis).

In this post, we consider a few of the approaches that we applied in determining that charity entrepreneurship is likely to be a high-impact career.

Approaches

Cost-effectiveness analysis for counterfactual impact

Expert consensus

Neglectedness

Historical track record (tractability)

Market inefficiencies (tractability)

Limiting factor analysis (tractability)

Comparative advantage

Cost-effectiveness analysis for counterfactual impact

A **cost-effectiveness analysis** outlines what sort of impact you get for your dollar. A more advanced counterfactually informed **impact evaluation** takes into account where the money and talent used in order to accomplish something might have gone instead. It factors in the losses that result from redirecting those resources toward something else, basically taking into account the opportunity cost of the project. Is it better to start a charity, for example, or to earn a six-figure income and donate 50k a year?

Cost-effectiveness analysis is most easily applied when comparing two similar options, such as distributing vaccinations to prevent disease-related mortality vs. distributing malaria nets to prevent disease-related mortality. This is because of the possibility of doing direct, **equally rigorous** apples-to-apples comparisons that consider the same factors.

When comparing apples-to-oranges options that are *not* as comparable to each other, you need to be more careful and skeptical of the final numbers. This is because it is difficult to apply equal rigor when comparing very different things. The choices we are currently evaluating in comparing the expected value of starting a charity against other career options, is sadly an apples-to-oranges comparison. However, a direct cost-effective analysis can still be a helpful tool for comparing impact – it should just not be used in isolation or taken literally.

One such analysis by Peter Hurford is posted on the [EA forum](#). Quite a few parameters go into this model, but a simplified version is that the average GiveWell top charity redirected \$8.3M per year from 2011–2015. This means that if a new charity has a 15% chance to redirect dollars in such a way that would do twice as much or more good as donating to GiveDirectly (the current cost-effectiveness bar to beat to become a top charity), that's an expected value equivalent impact of donating \$1.2M directly to GiveDirectly. If you split that impact between three senior staff, that comes out to about \$400k of GiveDirectly-equivalent-impact per person. You can play around with a more complicated version of the model in [Guesstimate](#), inputting your own parameters.

Charity Entrepreneurship's [cost-effectiveness analysis](#) is a bit more complex and based on the average of our predictions for 7 charities that went through the CE incubation program. Like the above you can modify the variables to see what the key parameters are and how the endline values change with different inputs. (Keep in mind that these results are generated based on the average of our predictions for 7 particular recommendations, so a single change of one charity's variables will only have 1/7 weight on the endline results).

CE's cost-effective analysis is anchored in estimates using historical accomplishments and observations from Fortify Health, Charity Science Health, and Charity Science Outreach (more established projects than our more recent charities, which have yet to reach their full potential). It also includes possible outcomes involving negative impact due to potentially wasting GiveWell funds (which otherwise would go to highly impactful charities), as well as a counterfactual/opportunity cost that assumes that each co-founder would have otherwise done something that achieved as much good as donating \$100k/year to GiveWell top charities (not an easy task). We tried to take the probability of success and the likely pool of donors for each individual charity we suggest into account, with the average result dictating the endline numbers. We also asked where the donors who supply our salaries and your seed grants would otherwise have donated, and subtracted the impact of those amounts.

As you may have noticed, cost-effectiveness analyses are difficult to get right, and reasonable people can disagree by an order of magnitude about expected value. For example, our analysis is quite sensitive to guesswork regarding the longevity, cost-effectiveness, and size of a successful charity, and Peter's analysis rests on assumptions about GiveWell funding landscapes that may not always apply. That is why we do not [take expected value estimates like this literally](#) – they are a tool to compare options, not necessarily a way to measure absolute value. That said, these

estimates do inform our endline view on the promisingness of CE as a career, relative to other careers.

The bottom line is that we think from this perspective that attempting to start a charity via our program does a huge amount more good for the world in expected value than the vast majority of people would realistically be able to accomplish by earning to give (i.e. trying to earn a high salary to donate a large amount to charity), as well as at most other careers. The number of great charities founded is often bottlenecked by the number of entrepreneurs who would be a great fit for the path. We have often seen great ideas that do not get founded due to no one being a great fit and confident enough to start the project.

Although this is a valuable perspective, it's far from the only angle you can consider when evaluating a choice as important as your career. Let's next take a look at what experts in the field think.

Expert views

Experts in the field often think you should start new charities. Elie Hassenfeld, co-founder and CEO of GiveWell, **wrote that GiveWell would like people to create new charities** working on **priority programs**, and would be willing to support the creation of these charities in the form of **incubation grants**. GiveWell co-founder Holden Karnofsky has also **spoken positively about starting organizations**. In conversation, GiveWell staff conveyed to us that another top charity could result in more dollars moved to top charities overall. More charities can translate to more capacity to efficiently use funds, and a greater variety of charities to fundraise for may allow for catering to broader preferences. Historically, this has been the case for GiveDirectly, which can quickly use large amounts of money and has a unique appeal to many donors.

80,000 hours has **identified** charity entrepreneurship as a potentially high-impact career that it classes as "sometimes recommended". 80,000 hours also suggests potential drawbacks of this career, observing that "the people who take this path seem highly able, suggesting it's very difficult," and it may take "several years to build up the necessary expertise". They are also "uncertain how easy it is to find and take these opportunities before they would be exploited by other nonprofits or governments." These are legitimate concerns and explain some of the reasons going through a handbook, course, or program can greatly improve your odds of having an impact.

Conversations with charity sector experts in the field (e.g. leaders of top recommended charities in the global poverty and animal space) also echoed the sentiment that a large amount of impact is possible, but that charity entrepreneurship is a challenging career path and not a good fit for everyone.

Finally, other prominent voices within the EA community have also said that founding new effective charities is a high-impact thing to do on the [EA Forum](#). And the wide range of support and interest the CE program has attracted is further evidence of this.

Neglectedness

There is a severe lack of effectiveness-minded organizations. When faced with the idea of starting a new organization, many people ask, “but aren't there already lots of effective organizations out there?” The truth is, effective organizations are rarer than you might think. David Anderson, Assistant Director at the Coalition for Evidence-Based Policy, writes:

“...what’s not often recognized is how rare—and therefore, valuable—such examples of proven effectiveness are. Their scarcity stems from two main factors: 1) the vast majority of social programs and services have not yet been rigorously evaluated, and 2) of those that have been rigorously evaluated, most (perhaps 75% or more), including those backed by expert opinion and less-rigorous studies, turn out to produce small or no effects, and, in some cases negative effects.”

The nonprofit world is incredibly small. There are [~167,000 registered charities in the UK](#). This may seem like a lot, but only 5% of them are large enough to employ any paid staff. Most of these charities are devoted to local issues. Even within the category of international charities, most run interventions that are significantly [less cost-effective](#) than the ones effectiveness-minded donors are interested in. Out of [785 charities](#) that GiveWell considered, they’ve only recommended eight as top charities – and bear in mind that GiveWell generally only considers charities with some degree of promise in the first place. In fact, GiveWell has identified multiple priority programs without any promising charities that focus on them at all – especially not charities that are eager to share information and can provide evidence that their intervention is working.

You might look at a given intervention that we recommend and notice that other organizations claim to operate in the area. This does *not* mean that the field is overcrowded and that there is nothing left to do. When an investigation reveals

other organizations working on an issue, this might mean one or two people working part-time, or a program with many different arms, of which the issue at hand is a small branch that doesn't get much resources. It might even be an organization that genuinely is working on the issue, but not necessarily in a cost-effective or transparent manner, even if it is generally a well-run organization. Even **when comparing programs within a given intervention**, some programs are orders of magnitude more impactful or cost-effective than others. To give an extreme example, imagine if someone had said: "Charity Navigator already exists. Why do we need a GiveWell?" Clearly the quality and organizational focus of GiveWell is quite different from other charity evaluators, and GiveWell is likely one of the highest impact charities ever started.

It's also worth noting that it can be high impact to start multiple charities in similar areas but with somewhat different focuses or country targets. There is more than one deworming organization on GiveWell's top charity recommendations, and GiveWell has noted that they would be excited to see even more charities attempt bednet distribution, deworming programs, and cash transfers. **The existence of similar charities is not necessarily an argument for not creating a new one.** The world is a big place and there is room for many nonprofits in it.

There is lots to be done and not enough organizations doing it – especially not organizations that are provably effective according to the standards that we hope to help you **achieve**.

(Note: I've mostly talked about the global development space here, but these points are even more true for other cause areas. Both outside and inside of effective altruism, global development is one of the most well-developed cause areas as far as evidence-based, impact-focused organizations go. Other cause areas are more severely neglected in this respect, often lacking empirically focused organizations altogether.)

Historical track record:

New effectiveness-oriented organizations with inexperienced founders have had some success. There is a perception that founding high-impact charities is near impossible, or at least unlikely unless the founders have decades of experience and specific related credentials. However, if we look at the historical evidence, it seems that this perception is misguided.

The following charities were founded by people with **fewer than four years' experience** in a closely related field and are well respected by the effective altruism

community (one of the most skeptical communities when it comes to charities). This is not an exhaustive list.

- GiveWell
- Against Malaria Foundation (GiveWell top charity)
- Charity Science Health (GiveWell incubation grant recipient, co-founder overlap with CE)
- Fortify Health (GiveWell incubation grant recipient, mentored by CE)
- Giving What We Can (4,536 members, donated millions to effective charities)
- 80,000 Hours
- Animal Charity Evaluators
- The Good Food Institute

What do these organizations have in common? All were founded by people who explicitly decided to do the *most* good. To give one example, the Against Malaria Foundation **began** when founder Rob Mather started a charitable swim to raise money for a burn victim. Realizing that he had a talent for fundraising, Rob asked himself which cause would do the *most* good to direct money toward. This simple intention of aiming from the start to do the most good, and spending some effort to figure out what that might be, was likely the key difference that set Rob apart from many other charity entrepreneurs.

The following charities were founded by people who may have had more than four years of experience, but were also either informed by or aligned with effective altruism principles from the start. This is also not an exhaustive list, just a sampling.

- Evidence Action (GiveWell top charity)
- New Incentives (GiveWell incubation grant recipient, mentored CE)
- The Humane League
- Machine Intelligence Research Institute

These are the successes. What about the failures? There seem to have been a comparatively low number of failures among people who explicitly aim to do the most good with their careers. When you systematically aim toward the high standard of being one of the world's top interventions, even your failures may ultimately have net positive consequences. Charity Science Outreach might reasonably be called a **failure**, but the failure created helpful information about

fundraising, and **netted \$220,000 CAD** over 2.5 years using two staff. While that admittedly may not be as high impact if you consider the counterfactuals of what the staff could have worked on instead, even if just earning to give, it was still net-positive. Not too bad for a failed experiment!

Charities conceived by the Charity Entrepreneurship research process in particular also have a strong track record. Both Fortify Health and Charity Science Health have received multiple GiveWell Incubation grants, and GiveWell has **estimated** that Fortify Health has a 25% chance of becoming a GiveWell top charity by 2022.

Based on this evidence, it seems possible that substantial prior experience and credentials are not necessary in order to start a highly effective charity.

Exploiting market inefficiencies

Charities are not efficient markets for doing the most good, so there is a lot of **low-hanging fruit**. What accounts for the unusual success of people who are aligned with or guided by effective principles? In the for-profit world, economists talk about “**efficient markets**”. This is basically the idea that because everyone is incentivized to make a profit, whenever an easy, low-risk way to make money appears, the opportunity is immediately snatched up and ceases to become easy and low risk.

It would be wonderful if people immediately snatched up opportunities to do good in the same way. Sadly, it's fairly rare for a charity to be established with the explicit goal of doing the most good. Some charities are started because the founder had a passion, an exciting idea, or a pet theory. Others are started because the founder was personally affected by an issue, or because the founder happened to have a personal fit in terms of skills, credentials, and connections. Most **donors are not discerning** about what their donation does, so there is **little incentive** for organizations to provide specific measurable outcomes for the donor's money. The prevailing culture is one in which criticism is frowned upon and a lack of transparency is the norm.

Very few people start a charity with the mindset of doing the most good within an empirically-backed consequentialist framework and begin with a systematic search of scientific literature and expert opinion. Few have the support of a community that encourages transparency and promotes more rigorous criticism as a cultural norm. These strengths mean that there may be low-hanging fruit in the

high-impact space that can be picked by even an inexperienced person, so long as they have the right priorities and mindset.

Limiting factor analysis

Charity entrepreneurs are well-positioned to address limiting factors. Before trying to solve a problem it helps to be aware of the **limiting factors** you will run into. For example, it may not help to fund **surgeries in the developing world** if the first bottleneck turns out to be the number of surgeons. Of course, you may be able to hire more surgeons with more money – but you would quickly hit a point where this is no longer cost-effective.

So what are the limiting factors for charity entrepreneurship? It is common for new charity founders to be intimidated by the prospect of fundraising, and most people immediately assume that the limiting factor must be funding. However, this is not necessarily true. **GiveWell's interest** and possible funding from the Open Philanthropy Project, **Good Ventures**, the **Global Innovation Fund**, **Evidence Action**, the **Gates Foundation**, and the **Children's Investment Fund Foundation** make us optimistic that global health projects will often receive adequate funding in early and middle stages. **The entrance of Open Philanthropy to the animal welfare space** means that animal charities that can demonstrate alignment with effective altruist principles currently exist in a very funding-rich environment and will have no trouble raising funds at all.

While the true **limiting factor depends on your cause area** and your specific project, it is often not as simple as “not enough funding”. There is scope for your work to unlock new sources of funding, and to cause existing funding to be used more efficiently than it currently is. If you follow all the principles of effective charity, we suspect that you will often find raising enough funds to be easier than you think.

Our research program aims to identify interventions that we think a determined person armed with the right tools can plausibly make progress on. For these interventions and many others, we think the true limiting factor is the number of talented persons such as yourselves who choose to work on highly effective interventions. A talented charity entrepreneur can increase capacity in any cause area, develop a new cause area, or improve the evidence-base in fields that have not attracted much professional scientific attention.

Summary

We've evaluated the notion of starting new charities (especially via the Charity Entrepreneurship program) using a variety of tools, integrating multiple perspectives to form our view.

- Explicit numerical cost-effectiveness analysis suggests that attempting to start a highly effective charity is an efficient use of funding.
- Experts in the field think it's likely to be a good idea for people who are the right fit.
- In the current state of the field, there is a lot of good work to be done, and unfortunately no good incentive structure is in place to push people to do it.
- Both Charity Entrepreneurship specifically and effective altruists in general have a good historical track record of succeeding in this endeavor.
- The limiting factor to success often seems to be the pairing of a talented co-founder and a strong intervention, rather than extreme funding shortage or severe implementation difficulties.

Taking all these factors into account, we think that trying to start a charity that implements an intervention chosen via a systematic process to maximize impact is an extremely valuable and cost-effective activity.

Reading list

Resources by Charity Entrepreneurship

- [Why founding charities is one of the highest impact things one can do](#)
- [Why EAs in particular are good people to start charities](#)
- [What is the expected value of creating a GiveWell top charity?](#)
- [Scaling Down Charity Science Outreach](#)
- [The Strengths and Weaknesses of Different Cause Areas and Charity Ideas](#)
- [Why We Look at the Limiting Factor Instead of the Problem Scale](#)

External resources

- [Sequence thinking vs. cluster thinking](#)
- [Why we can't take expected value estimates literally \(even when they're unbiased\)](#)
- [GiveWell Incubation Grants](#)
- [GiveWell Priority Programs](#)
- [GiveWell Charities we'd like to see](#)

- Founding effective nonprofits (international development) – 80k Career review
- Some causes are better than others – 80k
- History of Against Malaria Foundation
- An impressive failure – CFAR on Charity Science Outreach
- Fortify Health — General Support (2019) GiveWell grant
- Broad market efficiency
- How Efficient is the Charitable Market?
- Obstacles to “giving as consumption”

2. Impacts on Your Future Ability to Do Good

Author:

Joey Savoie,
Co-founder & Director of Strategy

The previous article explored the expected value of founding an impactful charity from a range of perspectives, focusing on the largest and most direct impacts. However, these impacts are far from the only ones. By founding a new charity, you can impact your and your team's future ability to do good, and influence the charitable movement you support. We consider this collection of benefits to be the **non-direct impact** of charity entrepreneurship as a career. This article is about the impact of charity entrepreneurship on your **future ability to do good**.

Doing good year to year is important, but many are concerned with making sure they can continue to have a large impact throughout their whole life. When considering this factor in a career context, take into account not just the direct impact of a job, but also its impact on your future. Charity entrepreneurship as a career path builds skills, credentials, background, knowledge, and therefore ability to do good in the future.

Credential building

Building up credentials, particularly if you are young, can be a highly impactful use of time. Credentials can be seen as **tools that get you a job**. Many employers like to see entrepreneurship on a résumé as it suggests an ability to work without a high level of management. Founders often take CEO or managerial roles, and the greater responsibilities that come with them, much earlier in their career than had they joined a larger organization.

Being in a senior position at a young organization often leads to a strong résumé for more senior positions in larger organizations, allowing you to move from one management role to another. In such roles, you are able to steer established NGOs and donors on a path of evidence-based development, thereby touching many lives.

People who have founded organizations tend to have impressive answers to common interview questions, like “how many direct reports have you managed?”.

Skill building

This refers to the abilities you acquire from working a job that could be cross-applicable to other careers – particularly careers that are also impactful. If credentials are the tools that get you the job, skills are the tools that will help you succeed in that job.

Building skills is even more important than building credentials because becoming highly skilled enables you to work your way into a job through hard work and ability (e.g. being hired from an internship, being headhunted). A strong résumé alone cannot keep you in a job you are bad at for long.

While building skills very often enhances your résumé career capital (e.g. starting a charity builds a ton of skills, and that gets recognized on a CV), it does not always work the other way around. You could get a great looking education that is useless in most jobs and builds very few skills. Although charity entrepreneurship builds both skills and credentials, it builds skills quicker than credentials.

Due to the wide range of challenging tasks, entrepreneurs learn a lot every day. By the time you’ve mastered a skill, you often end up hiring an employee to do that aspect of the job while you go on to learn about a different area. Picking up and mastering new skills quickly is a highly valued skill – one that is trained through charity entrepreneurship.

Attributable impact

It’s easy to imagine that you might miss out on a highly impactful job due to a gap in credentials or skills. However, one of the biggest ways people fail to have a long-term impact is by working on something that is not impactful without knowing it.

The nonprofit world is hard to quantify. Knowing how much good you have done is much more difficult than measuring how much profit you have made. Evaluating a charitable organization is a major challenge, often requiring a large external team such as GiveWell to consider all the details and calculations. Calculating personal

impact is often even more challenging, as it requires a sense both of your organizational impact and of your personal impact on the organization. Counterfactuals also affect personal job calculations strongly. For example, becoming a doctor seems like a great way to improve lives until you consider that had you not become a doctor, **someone else would have taken your spot in medical school and become a doctor** instead.

When calculating the impact of charity entrepreneurship, **you can simply calculate the impact of the organization as a whole**, avoiding some of these counterfactual replaceability concerns. In 2019, we were bottlenecked by the number of sufficiently talented applicants. Had we found one more person who cleared our bar for admission, it's likely that an additional charity would have been founded, indicating that charity founders going through our program have **low replaceability**.

Primarily needing to take into account organizational impact alone, instead of both organizational impact and personal impact within that organization, is rare in the charity world. By starting a charity, you may get a more accurate sense of your impact than most people can. Organizations are also far more likely to be externally evaluated than an individual position within an organization, giving unbiased data on the founders' impact.

Passive impact

When considering long-term impact, one of the unique factors to consider when starting a long-term project (charitable or otherwise) is the importance of setting up the project to continue after you have moved on. Most impact evaluations are generally done year to year and do not include ongoing long-term costs. But much of the benefit of starting an organization compared to, say, **earning to give**, is that **you can step back and the organization can keep creating impact**.

Some founders stay on a successful project for the rest of their lives, but many move on to other high-impact projects. If a founder starts a successful charity and then moves on after 5 years, finding a director to replace them who runs the charity equally well, the organization can continue to have an impact in addition to the founder's next career. In the for-profit world, this would be called a passive income. A company could continue to generate revenue with relatively little input from its founder. In the nonprofit world the same concept can apply, and a founder can have passive impact from previously-started, still-running charitable projects. Serial

charity founding can lead to multiple high-impact projects running, all thanks to a persistent co-founder. This can allow future impact to be doubled or more.

Personal value drift

Another common way people fail to do good in the long term is **a change in values or life circumstances**, such that impact is not as important as it once was. A person might collect a lot of career capital when they are young, with intentions of doing something altruistic and high impact later in life, but find that as they grow older priorities and responsibilities and interests beyond doing good begin to take precedence. The end result of this is often that less good gets done. Value drift describes how over time, people become less motivated to do altruistic things. (This is not to be confused with changing cause areas or switching to a different method of doing good.)

There is a strong precedent for value drift among other related concepts, both ethical things (such as being vegetarian) and things that generally take willpower (such as exercising). Having an impact through your **career is no exception** to this rule. There is convincing data that **value drift can be a major source of lost impact across your life**.

Thankfully, there are a large number of things you can do to **mitigate the risk of value drift** and increase your chances of having an impact long term. Many of these strategies involve surrounding yourself with impact-focused individuals. Having a job in the charity sector keeps you engaged with doing good, and charity entrepreneurship is no exception. Maintaining a good workplace culture with a supportive, impact-focused peer group can be a way to greatly increase your chance of focusing on impact in the future.

Overall

The decision to make the leap and start a charity could have a massive effect on your ability to do good throughout your life. It will teach you many skills; build strong credentials; allow you to make a large, clear, and measurable positive difference to the lives of others; and keep you engaged with a community of extremely talented people who are dedicated to doing good. This handbook aims to help make this process as successful as possible.

3. How Your Team Benefits From Charity Entrepreneurship

Author:

Joey Savoie,
Co-founder & Director of Strategy

Our default when considering impact is to consider our individual output – our personal work hours, money moved, and impact achieved. However, a significant flow-through effect of founding a new organization is that it can enable others around you to have a higher impact than they would otherwise.

Job opportunities

Many impact-focused individuals and effective altruists want an opportunity to work in an effectiveness-minded organization. It's easy to see why, as these jobs provide a strong opportunity to build skills while having a direct impact. They also provide a like-minded community, which is one of the biggest predictors of job satisfaction. However, these jobs are very **limited and difficult to get**, leading to frustration among job seekers. Founding new organizations in promising areas with an effectiveness mindset is a way to create more impactful positions that job seekers can apply for and use to build their own skills, career, and impact. Many effective altruists and charity workers who later went on to do highly impactful activities talk about their earlier jobs building their capacity for these roles.

Jobs also have more benefits than just the direct skills built on site: they can create strong bonds that outlast specific positions but lead to impact down the road. For example, some early contacts I met at EA jobs are still close friends. Their presence in my life greatly increases the scope of my impact and decreases my odds of **value drift**. A final benefit of creating job opportunities is that by increasing the variety of available jobs, you can increase the range and diversity of people who can find a high-impact way of contributing to the world, leading to a more diverse movement.

Staff training

When school ends, your real education begins. Learning on the job is an essential way to build skills. When you found a new organization, you create not only jobs but also pathways to more quickly building up employee skills. Some skills are highly difficult to learn outside a job that encourages the application of them consistently. When a movement is missing skills (and **every movement misses some skills**) one of the ways to build capacity is to have training built into impactful organizations. Impact-focused organizations often realize a portion of their impact from building the capacity of people within their organization, even if they do not stay indefinitely. With strong staff training, several employees can take on far more ambitious projects than they would have been able to otherwise. This can be particularly important if your charity focus is on an area that is underutilized in your broader cause area (e.g. measurement and evaluation in a cause area that has historically been weak at this).

Building a Schelling point

Often when a new charity is founded in an area, there is a boom of interest. Volunteers are keen to contribute, donors to support it, and other projects to link up. Having a central point of contact for a specific issue or concern, or a default place that someone with interest in an area can go to (a **Schelling point**) can make the difference between key members of a community connecting or not. Some nonprofit organizations state that their main benefit is providing this connection between different members of a community. For example, our organization, Charity Entrepreneurship, has become a Schelling point for people interested in founding new charities, the default place to contact and speak about related ideas. We have been able to connect people who otherwise would have had a challenging time connecting (e.g. two founders both interested in the same idea but waiting for the right co-founder). Many organizations in neglected areas can have the same effect.

4. The Impact of Charity Entrepreneurship on the NGO Sector

Author:

Joey Savoie,
Co-founder & Director of Strategy

Founding a new organization has a wide range of effects on the charitable sector as a whole. If a charity is good, it can raise standards in a field, spread or stabilize key concepts, build the EA movement, or encourage the allocation of more resources to an area. There are hundreds of effects that a new charity can have on the sector as a whole, both positive and negative. This post covers a few of the most important positive ones.

Raising the bar

The presence of an effectiveness-focused charity can raise the bar for a whole cause area. It allows donors and staff to have higher, but still realistic, expectations for other charities. In areas that are not very effectiveness-focused, donors will often not even know what an impact-focused charity would look like. Yet after they see one, they can become a strong advocate for expansion of the principles.

This higher bar can cause multiple charities to strive to be more effective – for instance, GiveWell reinforces its top charities' focus on cost-effectiveness and encourages every organization to strive for higher impact.

Spreading concepts

Sometimes the biggest impact an organization has is in its concepts, not in its direct impact. Many concepts and ideas from the effective altruism (EA) universe have not spread far or are limited to a particular cause area even within EA. Often, a highly impactful charity can be founded by someone who brings knowledge and a way of thinking that is uncommon in that area. When a new charity enters the space with

concepts like a focus on cost-effectiveness or counterfactuals, these models can spread to nonprofit partners that otherwise would never have heard of them.

Some of these concepts are easy to understand and can lead to a large difference in the way charities are run. Charity-to-charity transmission can happen through staff who learn lessons and then move on to other organizations, and by working in partnership on projects. **Charity Entrepreneurship incubatees** were inspired by other great charities and will be the inspiration for future charities yet to be founded. Good habits and ideas can be passed on, as can bad ones – which increases the importance of running an impactful organization.

Influencing donors

Donors are affected by strong charities as well. A donor might like that you are so transparent, and then make it a requirement for other charities they fund. Frequently, funders and others are affected by concepts they learned in one field and can apply in others. For example, over time The Gates Foundation became a supporter of open science, and now all research they fund is published openly. This change is arguably more impactful than any grant the foundation gave in the open science field, but was likely caused by their initial interest and investment in the space.

Impact on the effective altruism movement

Sometimes, a charity or individual goes further than just learning a few concepts, and can become a member of the growing effective altruism (EA) movement. The effective altruism movement is still largely populated by science and philosophy majors in the USA and UK. By having more organizations and jobs in EA-minded organizations, the movement can slowly expand and diversify.

Many of those who would most benefit from effective altruism concepts are those already implementing interventions in the field. Field staff are usually not hired with a strong effective altruism background, but get acquainted with the concepts through a daily connection to an organization they feel is doing good work. Similarly, funders often hear about effective altruism via connections to a single cause area or charity. There are many areas where there could be an effectiveness-minded donor who has not gotten much exposure to effective altruism concepts simply because the movement has little presence in this area.

Environmentalism, family planning, and mental health have many donors who could be interested in effective altruist concepts but have not heard of them.

Bringing in people with these backgrounds also has large benefits for the EA movement. Historically, there have been concerns that there is a lack of practical knowledge and understanding of the cause areas. Starting something as practical and direct as an impactful charity can show that the EA movement practices what it preaches and tests its abstract considerations and philosophy in the real world, where feedback loops are quicker and clearer. These feedback loops can provide key information that a purely theoretical movement would not be able to access.

New organizations can also provide a stabilizing force to the EA movement. Organizations tend to outlive social movements. Nonprofits tend to be slower to change or fade than the movements they are part of. As a result, they can give a movement multiple points of failure. It can also take the pressure off a single organization – if a movement is supported or endorsed by many nonprofits, a single mistake made by one of them will result in fewer ramifications to the movement as a whole.

Drawing resources into an area

If you're a strong charity in an area, people will have more positive thoughts about that area in general. For example, if you start a strong charity in tobacco taxation, more EAs will be interested in the area, and more donors will consider it as a possible field and support the area as a whole.

This can also work on a broader scale. If you make a strong case to funders that global health is a promising cause area, you can also direct interest from one very broad area to another; for example, from microloans to global health, or from environmentalism to animal welfare. If funders feel that your area is impactful as a whole, they might donate not only to your charity but also to other charities in your cause area. When you run an organization, you will have supporters who back you for purely impact-based reasons, but you will also have supporters who are interested in a different reason.

Sometimes an area looks promising, but a single limiting factor is preventing it from moving forward. If a new organization can reduce this limiting factor, it can open up large pathways to impact. For example, some areas have very large and interested funders (e.g. tobacco taxation, family planning) or have a new, much

larger source of funding (e.g. natural disasters, the animal movement). Donors interested in these areas can struggle to find **room for more funding** among impactful charities. Sometimes, the only areas left to fund are marginal improvements on a good charity (after they have accomplished 90% of their impact/effort) or large improvements on a disappointingly weak charity. Both are not attractive cases for donors to support. One new impactful charity, however, can create a considerable amount of “room for funding”.

Government interest and uptake is another possible limiting factor that charities can help with. Governments tend to be very slow-moving, and are often more likely to take on a project after it has been proven to work in their specific context. A strong, measurement-focused charity can give a government the confidence to run a national or state-level program.

Not only do charities often serve as proof of concept, but they also create some norms and methodologies around implementation that can later be used by a government. Of course, often the charity will work directly with the government to apply for a program after the charity has run it on a smaller scale. This process can allow an area to receive a huge amount of resources that the government would otherwise have spent elsewhere.

5. The Benefits on Top of Changing the World: Four Reasons You Should Consider Charity Entrepreneurship (Besides Impact)

Author:

Patrick Stadler,
Director of Communications & Charity Mentor

Helping hundreds of thousands of beneficiaries with an evidence-based and cost-effective program: that's your ultimate goal as a charity entrepreneur. It's not a coincidence that many benefits of becoming a charity entrepreneur are related to impact (see our article on the [impact of CE](#)). Yet the advantages of starting your effective nonprofit go beyond impact. As a founder, you will grow in various ways.

Here are four advantages of becoming a charity entrepreneur besides impact:

1. Skill building: from rookie to start-up leader
2. Satisfaction: find excitement in autonomy
3. Personal growth: a guarantee to be challenged
4. Exposure: collaborate with your heroes

Skill building: from rookie to start-up leader

Every day a new challenge, often one you have never experienced. The typical day of an entrepreneur requires rapid and constant skill-building. Preparing a budget in the morning, conducting a fundraising interview at noon, onboarding new staff in the afternoon, and working on your one-year strategy in the evening. As a charity entrepreneur, you are required to grow various skill sets at the same time and become a jack of all trades. In no time you will have a basic understanding of nearly every aspect of running an organization, from high-level decision-making to the technical intricacies of your IT systems. Such a broad skill set is highly transferable to any organization in the nonprofit or for-profit space.

With your ability to both think strategically and implement rapidly you will stand out. Most people excel at either the strategic or practical level. Think of all the highly intelligent individuals who love discussing research and theory but have no

track record of accomplishing something in the physical world. As a successful charity entrepreneur, you have proven an ability to connect the strategic and practical levels, combining the best of smart thinking with getting sh... done.

Running a start-up, you love being involved in various tasks. Soon, however, you realize that you need to leverage your time to be successful since your personal work hours are limited. So you start building your organization and grow into the role of a leader and manager. Through your teams, you multiply the impact you could achieve as an individual. While you might initially not be very good in this new role, you will eventually master the toolkits of task management, planning, and coaching – another highly sought-after skill set that will allow you to contribute to various high-impact organizations in senior positions.

Satisfaction: find excitement in autonomy

Working on a variety of tasks not only increases your skills – it is also exciting and fun. Yes, it can be satisfying to obtain expert knowledge in a particular subfield. For the typical entrepreneur, however, it quickly becomes dull. A variety of tasks and challenges stimulates a curious mind and contributes to job satisfaction.

By nature, the job of an entrepreneur comes with a high degree of autonomy, which is another key ingredient of satisfaction. Autonomy increases motivation while reducing mental strain at work. Decision-making autonomy has a particularly large effect in the literature.¹ There is no shortage of decisions for a charity entrepreneur: from high-level questions such as in which region to operate the program, down to questions about a particular recruitment process. As a charity entrepreneur, you have the final say and therefore the highest level of autonomy.

Autonomy translates to flexibility. As a charity entrepreneur, you decide when and where you work. Good luck explaining your 11 am to 11 pm work routine to your supervisor in your country's public service! As an entrepreneur, on the other hand, you are the one who sets your schedule. The same applies to your work location. While you want to be based as close to your beneficiaries as possible, you are generally free to decide where you work. Type away as you sit on your apartment's couch, and nobody will bother you. Similarly, you might be able to travel and work

¹ Muecke/Iseke (2019) 'How Does Job Autonomy Influence Job Performance? A Meta-analytic Test of Theoretical Mechanisms' at <https://journals.aom.org/doi/10.5465/AMBPP.2019.145> (Accessed 11/13/2019)

remotely. Depending on your location, you may benefit from low living expenses. It is not unheard of to take work calls over the internet next to a beach in Southeast Asia and be fully productive.

Meaning is a key pillar of life, as authors such as Frankl² have pointed out. Going beyond the self and helping others is a profound path toward meaning. Both ancient belief systems and modern psychology come to this conclusion. The mission and impact of your charity are strong contributors to this which will be reflected in your job and life satisfaction. The effects can be felt in the present and in hindsight. When you look back at your life at 85 years old, you will be happy about the time spent as a charity entrepreneur, while you might regret the stint in a marketing agency promoting the latest diet snack.

Personal growth: a guarantee to be challenged

Beyond career capital and satisfaction, entrepreneurship triggers personal growth. Whether you have asked for it or not. Running a start-up charity, you face numerous challenges that push you beyond your comfort zone. Your intervention fails in a randomized controlled trial, your key partner is not delivering, a major donor pulls back: these are just some of the roadblocks you could face. Yet navigating them is a surefire way to learn more about your personality and weaknesses. As a result, you will grow not only in the narrow sense of your career, but as a human being.

As you can imagine, this process can be difficult, so be sure not to push yourself too hard. There's a fine balance between challenging yourself to grow and setting yourself up for a breakdown. Reach out for help from your co-founders, coaches, and trained professionals. Recognize that basically everyone feels similarly in your position. This is normal, although you might not see it, as there is a bias for celebrating success stories over realistically reporting the ups and downs of a start-up. As insiders of highly successful ventures can tell you, the inside picture always looks messier than the perception from the outside might imply. Nobody tells you how the (hopefully vegan) sausage is made.

What doesn't kill you makes you stronger. The saying also applies to entrepreneurship. As you master increasingly bigger challenges, you will notice how

² Frankl (1946) 'Man's Search for Meaning'

https://www.goodreads.com/book/show/4069.Man_s_Search_for_Meaning (Accessed 11/13/2019)

you feel more confident with high-stakes situations, public speeches and general uncertainty. You will also have a more realistic self-image as you better understand your strengths and weaknesses.

You might curse your choice of entrepreneurship at times, but you will grow.

Exposure: collaborate with your heroes

Running a start-up, you collaborate with your heroes in two ways. First, as you run the show you get to pick colleagues who share your values and work ethic. You have the privilege of working with committed individuals who you can constantly learn from, given that their talents will surpass yours in many areas.

Secondly, as a leader of an organization, you interact with the highest levels of established organizations and governments. How many young recent graduates sign a deal with a Ministry of Health or get to speak at a global conference? As a charity entrepreneur, you will gain exposure, build strong networks, and, most importantly, be able to work with super smart and nice people inside and outside of your start-up.

In sum, opt for charity entrepreneurship if you would like to have an impact. And you will magically get a full package of goodies that goes beyond impact: solid skill-building from rookie to start-up leader, satisfaction through autonomy, personal growth by mastering challenges, and the ability to collaborate with your heroes.

II. Disadvantages

1. Three Years Later

Author:

Joey Savoie,
Co-founder & Director of Strategy

Three years of hard work, long nights, stressful decisions, and emotional investment. The result of that work would all come down to this.

John and Alice had been working on a new charity start-up for years and had done a good job. Their idea had been picked via an intensive research process. They had piloted a small scale version of their program, which showed promising results. Their execution was best in the field. But now they were waiting nervously for an email.

This was not their first time waiting for an email that would decide the fate of their charity. From funder emails to critical advisor feedback, they had dealt with many tough emails – but none quite like this one. The email they were waiting for was the results of the high-quality randomized controlled trial (RCT) run on their project. They had some promising evidence from their pilot and other studies in different contexts, but there was still a solid chance this RCT would show minimal or negative effects and their project would not be seen as impactful, requiring a large pivot or shutdown. Good or bad, the results of this RCT would determine the fate of their charity and suggest the historical impact of their work.

They both were startled by the sound of an email arriving at their inbox. It was unpoetically named “RCT – results”. Alice opened the email and quickly read the contents. Her heart dropped. The results were bad – quite bad if the summary at the top of the email was correct. The study had found the program to have no significant effect.

It was hard to believe given all the hard work they had put in. They had done everything right and still lost. She looked up at John who was reading the results himself in disbelief. Trying to process that their charity was not a success, but a failure.

Charity entrepreneurship is a high-risk career. Even if you do everything right, new data can come in that shows your intervention is not effective. A government can change and the doors that were once promising can all immediately shut. A seemingly small argument with a co-founder can build until both people resent each other. Even for the best charities, there is an over 50% chance they end up having minor or no impact.

This harsh reality is hard to grasp in the abstract. Imagine flipping a coin at the start of your charity, and that if the coin comes up tails your charity shuts down as a failed project. Information can be gained from failed projects and skills can still be built; however, the psychological experience of failure is only minorly mitigated by these benefits.

This section talks about some of the disadvantages of working in charity entrepreneurship, including the possibility of higher-impact alternatives and the high chance of failure.

2. What Would Have Happened Otherwise?

Author:

Joey Savoie,
Co-founder & Director of Strategy

Note: “What would have happened” is also known as the counterfactual impact. For more details on this concept as an epistemological approach, turn to our article on [counterfactuals](#).

Charity entrepreneurship as a career path is one of the oldest and most commonsense ways to make a difference. In the previous chapter, we discussed the many positive impacts of charity entrepreneurship as a career path. Now we turn to some of the potential downsides, starting with **what would have happened otherwise**.

If minimal progress is being made on a problem despite for-profit and governmental efforts, this could be a reason to found a charity. However, the picture complicates when you take into account other charities and funding cannibalism. More importantly, the bar for charity entrepreneurship being the career to pursue is not just that it has an impact, but that it has a **higher impact than other careers** possible for an equally talented individual.

Other charities

Every charity requires staff, funding, and space that another organization could use. For this reason, one of many people’s biggest concerns about starting a charity is that the world is flooded with charities. Indeed, there are **167,000 charities in the United Kingdom alone** – that seems like more than enough charities. But the number shrinks when we look more closely at the data. For example, there are 5.7 million businesses registered in the UK – 34 for every charity – and few would say that there is no space for new businesses. Another important factor to look at is size. Turns out when you look at the data that 95% of charities are too small to employ a single paid staff. Half have a yearly budget of less than £2,000 (roughly 2,500 USD in 2020). These organizations are very small in scope, and in some cases completely inactive.

Another strong factor in the charity world is the huge amount of charities focused locally. Only **around 10% of charities** in the UK operate internationally. Some might think that surely most charities registered in the UK work in the UK, since if they worked elsewhere they would register in their country of operations. But in most cases, charities working internationally will register both in their country of operations and in the country from which they expect to raise funds (most often the UK or US).

Overall, this leaves a much smaller number of charities – about 800. Say each of these charities were focused on health care, and were the size of a standard UK hospital. That could be 800 hospitals across a given country. This could do a ton of good, but seems small compared to the 1,257 hospitals in the UK alone. Of course, these calculations are rough and some charities will be considerably larger than a hospital. But a quick look at the financial data suggests this might even be generous, with **the whole UK charity sector spending about £48 billion (less than Tesco, a single chain grocery store)** compared to the NHS' £129 billion. And keep in mind that £48 billion is the spending for the whole sector, not for international charities.

Compared to the scope of the problems they tackle (e.g. disease, war, poverty, environmental, and animal issues), the number of charities is pretty small. However, that still does not suggest that founding a charity is good. Maybe there is space for new charities, and issues that are uncovered. But if funding is fixed, if people donate the same amount – then a new charity is simply taking funding from another charity.

Money

Competition for funding within the sector means that an effective charity must be **cost-effective**. **Different sources** report somewhat different percentages of income donated to charity, but fairly consistently it looks as though individuals donate the **same amount to charity**. This leads to the concern of **funding cannibalism**. If a new charity is started, doesn't it mainly take funding from other charities to do its work? In short, yes. This has some interesting implications. It means that starting a low-impact charity draws funding away from other charities, so that less good is accomplished in the world. On the other hand, the impact of starting an unusually high-impact charity is the difference between that charity and where the dollars would have otherwise been spent.

For example, say that in a given sector of the charitable world, \$10,000 in donations can save a life. For a new charity that is less cost-effective than average – where

saving a life requires, say, \$20,000 in donations – successful fundraising could be net negative. This is because for each life this charity saves, the average charity could save two lives. The logic also works the other way. If a new charity could save a life for \$5,000 compared to the industry average of \$10,000, each time it saved a life the funding otherwise would have only saved half a life.

The implication is that to do good, **new charities need to be more cost-effective than the sector average**. The specific sectors will each have different bars for effectiveness. Some areas have worked in a numerical framework and generally are careful to measure outcome metrics such as lives saved. Global health stands out in this regard. But in many areas, charities are not focused on endline metrics and will have much lower average bars.

Time

Money is not the only resource your charity uses. Yours and your co-founder's time are valuable resources that could be used to do good for the world in other ways. The same is true for staff that you hire.

Similar arguments as in the case of money apply for the staff you hire. Most of them will work in the same area but for a different charity. If you are more cost-effective than the average other charity in the area, this will be less of a concern. If you are less cost-effective, this could majorly affect your impact. To calculate yours and your co-founder's impact you can consider your plan B if charity entrepreneurship did not work.

3. Six Ways Your Charity Start-up Might Fail, and How to Prevent That

Author:

Patrick Stadler,
Director of Communications & Charity Mentor

Fast forward: Your charity start-up just hit the ground running. You are passionate about your cause, you just gave a big talk at a conference, and you've set up a shiny website. All is looking great. Yet your organization might still fail: it might collapse or not have any impact. Here are six ways a charity start-up might fail – and how to prevent this scenario.

1. You picked an ineffective intervention.
2. You chose the wrong co-founder or partners.
3. You don't understand your beneficiaries.
4. You prioritize research instead of focused implementation.
5. Your fundraising is not working.
6. You transition leadership too quickly.

You picked an ineffective intervention

What can go wrong when you select an innovative idea and a cause area you feel strongly about? Potentially, a lot. Innovation and personal passion for your idea are weak indicators of your intervention's suitability. First, your personal fondness for an idea might make you overlook its weaknesses. Second, innovative ideas are often overrated. There are dozens to hundreds of proven, evidence-based, and cost-effective interventions that need to be scaled. And they are just waiting to be implemented, sitting idly in the databases of applied research organizations such as J-PAL, 3ie, GiveWell, or Charity Entrepreneurship. Picking one of these sets you up on a path of impact.

Remember that the intervention, not your implementation skills, is the principal limiting factor of your organization. If you select a poorly-researched and

expensive intervention, even having stellar implementation will not get you anywhere. Moreover, as a first-time founder, you can be fairly certain that your operational skills will not be outstanding, so a high-impact intervention area can compensate for that as well.

How can you avoid picking an ineffective intervention?

- Prioritize **cause areas in which you can assume high impact** based on current and future levels of suffering (e.g. animal welfare, poverty, mental health).
- Prioritize **developing countries**, where your funds go a long way and rapid population growth takes place.
- Build on the freely available research of J-PAL, 3ie, GiveWell, or Charity Entrepreneurship to **pick evidence-based interventions**.
- In terms of studies, prioritize **randomized controlled trials** over weaker study methods.
- Consider not only effect sizes but also **cost-effectiveness models**, as the cost per impact achieved is what you should optimize for. See GiveWell for sophisticated cost-effectiveness models.
- **Triangulate different sources** to arrive at conclusions. For instance, back up your findings from randomized controlled trials with an expert survey to confirm your reading of the evidence.
- **Be ready to switch interventions** if you learn that your path is not evidence based or cost-effective.

You chose the wrong co-founder or partners

The cliché is correct: **you are as good as your team**. Working with the right co-founder and partners is critical.

You are setting yourself up for failure if your whole organization is not aligned on values such as beneficiary focus, lean operations, and an evidence-based approach. Divergent working styles and preferences can also have a detrimental effect. For example, if you like to plan thoroughly and your partners prefer implementing at the last minute, a good working relationship will become tricky.

When choosing a co-founder or partners, try to list factors for a successful relationship in a spreadsheet to give you some degree of objectiveness and comprehensiveness. However, given the importance of the human element here,

also trust your gut as a suitable heuristic. A bad feeling about a relationship can be strong intuition not to proceed.

In terms of picking suitable co-founders and partners, look out for:

- **Complementary skill sets.** Don't clone yourself in terms of skill sets. Look for partners who have different strengths and weaknesses.
- Select for **entrepreneurial attributes** such as an inclination for proactiveness, a readiness to pivot, and comfort amidst uncertainty.
- **Exclude those who do not align on core values** such as following the evidence and opting for cost-effective interventions.
- Use trusted people in your network or an incubation program to **look for potential matches**; don't just expect the candidate to find you through a public job posting.
- Be thorough when doing **reference checks**. Go beyond standard questions by asking for specific responses (e.g. would you rate x in the top 10% of partners you have worked with?).
- Don't start the relationship as a fully committed, publicly announced collaboration, but **set up a trial** in the form of a consultancy or limited engagement – especially if you have certain doubts. You can always promote someone, but demoting is often a recipe for trouble.
- If you cannot find the above, **don't settle**. Continue without a co-founder or partner until you find the right match. Having the wrong partner is worse than being on your own.

You don't understand your beneficiaries

You know all about the effect sizes and p-values of the research papers in your field. You have a sophisticated cost-effectiveness model ready to be applied. Yet you have barely spoken to the beneficiaries of your intervention. This is an issue, as even the intervention that appears most cost-effective on paper is not worth much if it is not successfully applied in the field.

“Understand your customer” is the mantra in the business start-up world. It applies to the nonprofit domain as well, even though the feedback loops are less straightforward: you get your funding from donors, while you serve your beneficiaries. Hence, your beneficiaries are your real customers. Do you really understand their problems and preferences?

The reality is usually more complicated than what one would assume sitting at a faraway desk. There is a myriad of reasons, for instance, for which mothers do not access free vaccinations to protect their children from deadly diseases. While research goes a long way, conversations and testing in the field can be extremely informative and provide answers specific to your context.

To understand your beneficiaries, you could:

- Acknowledge that you cannot find all the answers in research papers and that you need to **get out there**.
- As a first step, **talk to experts** who have interacted a lot with your customers.
- Talk to the customers themselves in individual **interviews and focus groups**.
- **Travel and live in the country** where you plan on implementing your intervention.
- **Test your assumptions**; answers to surveys are much less valuable than actual tests (e.g. does someone really use product x that they claim in a survey they would surely use?). Tests should first focus on the beneficiary's understanding and then on the usage of your product, as the latter requires the former.
- Ideally, **embed yourselves in the realities and lives of your customers** for a longer period of time, e.g. work at a farm or farm supplier for a few weeks if you plan to roll out an intervention on animal welfare that should be adopted by farmers.

You prioritize research instead of focused implementation

Weeks or months into your implementation, your focus remains research. “Let’s just look at the evidence for another 3–6 months.” This intuition might be correct in some cases. However, for the large majority of action-focused charities, at some point you have to arrive at conclusions from your research and start the actual implementation. There will never be an optimal moment for this; just look at the standard request for further research at the end of each academic paper. Yet, unless you are a charity whose core objective is research, you have to go beyond papers sooner rather than later, and start implementation.

Prioritizing implementation over research is not sufficient. You also need to move quickly to keep donors and staff – and yourselves as founders – fully engaged. It is difficult to maintain momentum if an organization has not made major progress

over 6–12 months. The key here is a focused implementation that prioritizes the most important and urgent tasks.

The following can help you focus on implementation:

- Compare your intervention with successful charities such as GiveWell’s top charities, and you will see that there are many unanswered research questions in their domains as well (e.g. the academic debate around the impact of deworming).
- Define sufficient research conclusions in advance and stop doing research after that target is achieved.
- If new research questions arise, either postpone tackling them or allocate a specific amount of time in the spirit of **‘time capping’** (i.e. “let’s arrive at conclusions on these new papers over three days, and then move forward”).
- In terms of operations, **focus on the tasks with the highest impact and urgency**. Define these tasks on a weekly basis. At the end of every week, revisit your success in working on these high-priority tasks vs. lower-priority activities. If your key tasks involve having x conversations with potential donors and securing a Memorandum of Understanding with a state government, and you see that you have spent two days optimizing your logo and website, you have direct feedback to improve. See [How to Prioritize Your Time](#) by Y Combinator’s start-up school for a great summary of this approach.
- **Depending on your team’s natural tendency, either slow down or speed up.** In a setting of research-leaning founders, as for instance in effective altruism, you likely should speed up. Adopt an 80/20 attitude on most things, according to which you can get 80% of the results with 20% of the effort. Be more ambitious on your timelines for achieving major milestones. You can always change course if something is not working out, but without any action, you cannot improve.

Your fundraising is not working

As a nonprofit organization, you operate at a loss by default and therefore rely on constant fundraising. Established organizations might benefit from a one- to two-year runway, but, as a charity start-up, you might only have funding for 3–6 months.

Fundraising is a time- and energy-consuming task for founders, but there are a few recipes that can help:

- **Check whether you can lower your budget.** The easiest way to solve a fundraising problem might be to reduce your cost structure. Reasonable budgets might also increase your likelihood of fundraising success, as lean organizations are often appreciated. As a young organization, you often want to be especially lean as it allows you to be more flexible. For instance, switching programs is easier if you don't have whole teams of staff working on these programs. Of course, there are limits to frugality, and you should not value the founders' time so little that they have to take care of every nitty-gritty detail instead of focusing on the high-value tasks.
- **Get your plan straight.** It is much easier to get donors on board if you have clear and exciting deliverables for the next 6-12 months.
- **Get into sales mode.** Fundraising is, to a certain degree, a numbers game, and you will need to establish many relationships in order for some of them to result in major contributions. If you have spoken to two foundations in your cause area, why not ten? If you have applied to three social entrepreneurship competitions, why not triple that?
- **Be comfortable with rejection,** and don't take it personally. Read some articles by successful salespeople, or check out the failure CVs of accomplished researchers (e.g. [Haushofer](#)) to understand how many 'no's have paved the way for rockstars.
- **Set up an advisory board with experts and trusted operational advisors.** This not only helps you steer your organization, but also expands your fundraising network. Asking for advice is usually a good strategy to establish a stronger relationship that can eventually turn into fundraising opportunities. This applies to funders themselves, but also to associated experts who can open their vast networks of donor institutions.
- **Be transparent, patient, and consistent in your fundraising relationships.** Treat your fundraising relationships like good personal relationships. Be honest about your organization's progress, and keep in touch. Nobody wants to hear from you only when you need them to write a check.
- **Explore strategic fundraising opportunities.** A one-time donation is great, but a strategic fundraising partnership goes beyond this by setting up a regular funding stream over multiple years. Here, it is key to see the donor as a real partner who will gain something from working with you, from social recognition to tangible support (e.g. if you allow your partner to access your human or technical capital).

You transition leadership too quickly

Your charity has successfully tested and scaled an intervention. Expert staff are in place and structures built. You and your co-founder consider leaving the organization to start something new from scratch. This is a high-stakes decision. It can go right if you have built solid systems to steer the organization and replaced yourself with staff who excel in their respective domains. It can also go terribly wrong if you have no high-performing and value-aligned senior staff in place to take over.

Here are a few things to consider:

- **Starting a new organization usually takes longer than you expect.** Three to five years is a much more realistic timeframe than one to two years. And if you plan to run this for 6–12 months only, then you might want to reconsider.
- While your involvement in all aspects of the organization is key at the beginning, try to set up processes and staff that can replace you eventually, including in key domains. If you don't make this an active priority, it will likely not take place. The bias of a start-up founder is always to get things done her/himself; this is what you excel at and enjoy, so you will actively have to challenge your understanding of your role, and scale as a leader. **Envision a scenario where the organization runs just as well, or even better, without you.**
- A key ingredient to replacing yourself is **recruiting and training the right staff**. Fight the natural tendency to spend too little time on recruitment. Daily urgencies are successful at grabbing your attention but less important than dedicating your attention to long-term recruitment.
- **Plan your eventual handover and departure** months or years in advance. Set up specific protocols regarding what will be handed over to whom. A general culture of well-documented processes in your organization, from employee handbooks to step-by-step instructions, will help your successor hit the ground running.
- **Clearly define your future role in the organization** to allow your successors to succeed (e.g. part-time role, no role, board role, advisor to CEO role, etc.). You want to be available for some continued support while not restricting the responsibility of the new leaders.

The six ways your charity might fail are also six ways you can excel and stand out. Keep these points in mind as you tackle your specific challenges, and you will continue the journey toward evidence-based impact. And don't forget to have fun in the process!

4. Five Challenges a Charity Entrepreneur Faces – and How to Manage Them

Author:

Patrick Stadler,
Director of Communications & Charity Mentor

“Why take this risk?” Mentioning at a family gathering that you would like to become a charity entrepreneur might trigger looks of concern and opposition. What about the high failure rate of start-ups, the low entry salaries, the need to work around the clock? There are potential personal challenges related to charity entrepreneurship. Neglecting them would be dishonest. Yet often these challenges are either exaggerated or can be dealt with successfully.

Let’s go through them one by one:

1. Uncertainty and a high risk of failure
2. Financial limitations
3. Separation from friends and family
4. Work-life balance
5. Unglamorous work

Uncertainty and a high risk of failure

The cliché is correct: a high percentage of start-ups fail. In the case of charity start-ups, an impact evaluation might show negative results, your funders might move on to other cause areas, or your team might no longer be sufficiently motivated by the organization’s mission. Even successful start-ups involve a lot of uncertainty. It is often hard to predict how the organization will look just 6-12 months in the future. The risk of failure and the general sense of uncertainty have an effect on both a psychological and financial level.

While financial concerns often come up first, the underlying challenge about uncertainty is likely *psychological*. As human beings – even the more adventurous ones – we strive for a certain degree of stability in our relationships and tasks. A start-up provides anything but. Yet the concerns are manageable with a few simple strategies. For instance, **acknowledge the worst-case scenario of failure** in a visualization exercise inspired by stoicism. Even in this scenario, your efforts will not be in vain on a personal or societal level. The lessons learned will contribute to shared knowledge, while you will have personally benefited from the journey, since you worked on a meaningful project and gained valuable career capital.

The pressure generated by your venture will also trigger personal growth that you might not have achieved otherwise (see [The Benefits on Top of Changing the World](#)). Other simple techniques, such as sports and meditation, can help you take a step back when you feel overwhelmed by the uncertain outlook of your charity.

Financial limitations

Financial challenges related to a charity start-up are twofold. First, you might worry about what happens financially in the case of failure. Second, you might be anxious about lower salaries in the nonprofit space and limited financial freedom.

In dealing with the financial risk of failure, it might help to **perceive charity entrepreneurship as an investment** similar to university education. You will benefit in the long term – and, as a charity entrepreneur, you get paid a salary instead of being confronted with school fees. Further ways to alleviate your financial concerns include building a 6- to 12-month **financial buffer before starting**, and checking out the social security benefits of your country of origin (as a last resort). And don't forget that as a charity entrepreneur you are likely highly skilled, so it is realistic that you could find another well-paid job in the worst case scenario. (See also [What is a Reasonable Financial Safety Net for a Charity Entrepreneur](#).)

In terms of financial freedom, starting your charity will restrict you more than working for an established company. Maybe you will have to reduce your budget for vacations, the latest gadgets, or eating out. However, **becoming a charity entrepreneur does not equate to joining a monastery**. Your donors understand that you can only be productive if your living costs are covered. So you will not have to switch from smart to feature phone, stop going out with friends, or spend holidays on your balcony. Finally, as an entrepreneur running an early-stage organization, you are likely already cost-conscious, so cutting back a bit will feel natural to you.

Separation from friends and family

Your happiness depends to a large degree on your social network. Being away from friends and family can be challenging even for social individuals. Charity entrepreneurship will mean time away from your loved ones. You might have to set up a program in a rural Indian state or attend multiple donor conferences in a faraway global capital.

However, job mobility is not unique to charity entrepreneurship. It has become a common feature of modern working life. The autonomy you enjoy as an entrepreneur can even help you stay close to your social circles. Many tasks can be accomplished remotely. For deep concentration, a quiet remote setting is even better than a hectic field setting. Moreover, if you have a co-founder, you can split up the time spent in far-away locations among multiple people (see how such topics can be raised in a Founders' Agreement in [How to Strengthen Your Co-founder Relationship](#)).

Work-life balance

The first years as a charity entrepreneur might not be the best to start time-consuming new hobbies or plan extended holidays. Your work-life balance is likely to suffer as you start your venture from scratch, test it in the field, and recruit the first employees. 40 hours per week will likely not be sufficient. At the same time, the workload should also be put into perspective. For many charity entrepreneurs, the counterfactual would be climbing the career ladder in a highly demanding job such as management consulting. The work-life balance there is arguably worse.

As an entrepreneur, the distinction between work and life will become less clear. After all, you have started something you truly believe in, that you pursue with a high degree of autonomy. This is a different path to working for a conglomerate in an anonymous office building.

As your organization grows, delegation processes will also become more established, and you will not have to constantly extinguish fires on top of your leadership duties. Staff will take care of operations, and you will be able to focus on strategic questions. All of this reduces your workload and stress levels.

Finally, while working more than 40 hours per week is clearly beneficial in the beginning, the science of productivity [1] shows that more is not always better. As a smart leader, you can resist the urge to work constantly, and **consciously set windows for uninterrupted rest and recuperation**.

Unglamorous work

The movie version of your life as a charity entrepreneur might involve speaking at a prestigious conference and drafting research reports in a shiny coworking space in a global capital. The reality is often less glamorous. You might be sweating in a tiny hotel room in a foreign country as you try to fix your internet connection to update your organization's website.

Initially, your organization only has you and your co-founder(s) as employees. By definition, you will have to carry out all or most of the tasks. Moreover, the large majority of duties at an NGO is operational: revising budgets, talking to lawyers, drafting work contracts, and so on. This is particularly important to consider if you come from a research background and expect charity entrepreneurship to be an applied research fellowship, which it is not in most cases.

The upside of this breadth of tasks is **broad skill-building** and **career capital**. You also get to understand every process from scratch, a good remedy against the detached decision-making that hampers many larger organizations. Of course, you will be able to hire staff eventually so you can focus your time on more strategic responsibilities. In the meantime, you could **outsource** repetitive tasks or IT-related tasks to freelancers available through platforms such as [Upwork](#).

In sum, charity entrepreneurship brings challenges, there is no question about that. Yet there are many strategies to manage and alleviate them.

Summary

1. Uncertainty and a high risk of failure are intertwined with an entrepreneurial journey. It helps to visualize the fact that, even in the worst-case scenario of a shutdown, charity entrepreneurship is beneficial for the broader community (through lessons learned) and for yourself in terms of career capital and personal growth.

2. You will be more financially restricted than your colleagues working in the corporate world. However, you don't need to become a monk and give up your social life. Planning for a financial buffer and reminding yourself that you can always get a well-paid job are two additional ways to reduce anxiety about finances.
3. Being separated from family and friends is often unavoidable in a modern work setting. Yet as a charity entrepreneur, you can split field time with a co-founder, and get to work remotely.
4. Work-life balance as a charity entrepreneur might initially not be ideal, but is potentially better than in careers such as consulting. The additional work hours will be compensated with meaning and autonomy. You can also keep your workload in check by delegating and setting up standard processes.
5. Charity entrepreneurship involves a lot of unglamorous work. On the plus side, this allows you to understand your organization from scratch and build broad career capital. You also have the option to hand over tasks to freelancers initially, and staff members eventually.

Reading List

Internal resources

- [Four Reasons Why You Should Consider Charity Entrepreneurship Besides Impact](#)
- [What is a Reasonable Financial Safety Net for a Charity Entrepreneur](#)
- [Impact of Charity Entrepreneurship](#)
- [How to Strengthen Your Co-founder Relationship](#)

External resources (free)

- [Come to Your Terms](#): preparing for potential failure (Mindingourway)
- [Make Self-care Your Competitive Advantage](#) (First Round)
- [The Effects of Working Time on Productivity and Firm Performance: A Research Synthesis Paper](#)

External resources (paid)

- [Six Pillars of Self Esteem](#)
- [Grit](#)

III. Assessing fit & preparing

1. What Traits Make a Great Charity Entrepreneur?

Author:

Joey Savoie & Kat Woods
CE Co-founders

If after weighing the advantages and disadvantages you are convinced that charity entrepreneurship might be a high-impact and satisfying career, the next question you might ask yourself is whether charity entrepreneurship is a good personal fit for you.

Starting a charity involves a lot more risk, stress, open-ended tasks, and heavier responsibility than most other jobs. We've listed the key personality traits and skills we think are necessary for a successful charity entrepreneur. The good news is, your personality is dynamic. You can cultivate certain traits if you want to, so if you do not have one of these it's not set in stone.

Successful charity entrepreneurs are:

1. Resilient
2. Ambitiously altruistic
3. Results-oriented
4. Open-minded
5. Not afraid to admit mistakes
6. Self-motivated
7. Creative
8. Doing it for the right reasons

You're resilient

It goes by many names – grit, determination, resolve, resilience. The strength to keep trying no matter what obstacles crop up (and, believe me, there will be

obstacles) is absolutely critical if you are to succeed. Rome wasn't built overnight, and founding a top charity won't happen after a week of part-time work. **If you want to start a charity, you have to want it, even after your plans A through P have all failed.** You have to keep trying even after you've had a scathing review online, because once you're doing something big enough that strangers start to comment, there will always be somebody who doesn't like what you're doing.

Being resilient means being dedicated at the highest level. It does not mean getting stuck on a specific plan or idea without flexibility. It means aiming for the same goal over a long period of time. For example, if someone with grit were pursuing journalism, they wouldn't apply for the same position dozens of times and no others. They would apply for hundreds of jobs in journalism and if that didn't work, they might take an online course to build up their skills, or start reporting on events in a public forum to gather a following.

You're ambitiously altruistic

You want to help so many people over your lifetime that they wouldn't all be able to fit in a football stadium. You want to wake up knowing you are pushing the limits of what is possible. Most people want to make the world a better place, but the majority only go so far as to be nicer to those around them or put a little extra thought into a present. Those gestures are laudable, but running a great organization needs more **vision**, otherwise you'll be too tempted by easier ways to change the world.

You're results-oriented

You worry that a lot of charities, while well-intentioned, are misguided. They often accomplish nothing and sometimes even make the situation worse. Your response to knowing that things can go wrong is not to say that making a positive difference is impossible, but that you have to learn as much as you can about the situation before making a decision. You want the analytical, critical, rigorous, and empirical thinking found in the scientific sector to become the norm in the nonprofit sector, too. **The stakes are too high for decision-makers to value emotional appeals over evidence and results.**

You're open-minded

Good charity entrepreneurs always remain open to the possibility that any and all of their assumptions may be incorrect. If you are not open-minded enough to consider new evidence and update your beliefs and actions accordingly, you are much more likely to fail. As a small condolence, you probably won't realize that you've failed because, as Kathryn Schulz explains, "how you feel when you are wrong is identical to how you feel when you are correct". On the flip side, being open-minded will allow you to eventually outpace many of your peers because you will be able to steadily improve your model of the world.

So, even if you initially think that something is incorrect, approach it with an open mind. Have a 'scout mindset,' trying to understand situations and concepts as honestly and accurately as possible, even when this is inconvenient. Remember, **changing your mind is the ultimate victory**, because in those moments you are improving your model of the world. And how can we ever hope to fix a problem without understanding it?

You're not afraid to admit mistakes

Humans are world-class self-deceivers, commonly making excuses for bad decisions, rationalizing away negative outcomes, and constructing fantasies to replace unpalatable truths. Charity entrepreneurs need to be able to work hours on end for years and then admit that they made a mistake, or that the project isn't effective enough to continue pursuing. This requires a rare level of honesty. Many people end up burying their heads in the sand, rationalizing away negative information. Don't let pride come in the way of achieving the most positive impact possible.

You're self-motivated

As an entrepreneur, you have to convince many people that your idea is a great one, but **the first person you have to convince is yourself**. You have to be able to get yourself up in the morning without a boss breathing down your neck. You have to motivate yourself to do unpleasant but necessary tasks. It's difficult, and some people just can't get the work done without a push.

A good way to proxy this style of work is to take an online course. There are thousands that you can take for free at your own pace. With such a powerful resource publicly available, it is amazing that people pay such huge sums of money for a university degree, but it comes down to motivation. Most people cannot complete an online course by themselves without a teacher guiding them to the finish line. As an entrepreneur, you can set up a board or peer group to help you with this, but when it comes down to it, being able to motivate yourself is crucial.

You're creative

Of course, coming up with the initial concept takes creativity, but that's just the tip of the iceberg. Almost every day entrepreneurs need to devise and compare multiple solutions to any given problem. Great entrepreneurs aren't afraid to think outside the box, do things differently, and bend or break social norms. **Difficult problems require creative solutions.**

That being said, it is possible to have a strong team that can collaborate creatively to build a strong organization. Not all original thinking has to come from the top. We find that **encouraging board and staff members to contribute and collaborate freely is a great way to produce brilliant ideas.**

You're doing it for the right reasons

There are many reasons to start a charity, and not all of them are altruistic. Some people do it to impress others, to have adventures, or to feel good about themselves. If you let these kinds of motivations interfere with the ultimate goal of helping, people will be harmed. For example, if your primary motivation is the warm glow of assisting others, but you find out that instead of ministering to the ill, it's better to prevent the disease in the first place, you may choose to stick to ministering because prevention is not as emotionally rewarding as treatment. Many people will lose their lives because of your misguided motivations. Likewise with prestige: a desire to impress others may cloud your judgment. Sometimes the best thing for a charity is to give the credit to somebody else.

Summary

How important these are in order (including some traits not discussed in depth):

- Doing it for the right reasons
- Ambitiously altruistic

- Results-oriented
- Personal fit
- Proactive
- Resilient
- Having a high general ability
- Self-motivated
- Open-minded
- Not afraid to admit mistakes
- Organized
- Creative

2. Which Jobs Will Best Prepare You To Become A Charity Entrepreneur?

Author:

Joey Savoie,
Co-founder & Director of Strategy

The first thing to consider is whether you need a job to prepare yourself for charity entrepreneurship. A number of people have been surprised at the impact they can have relatively early in their career. In fact, we generally think that having the right goals and personality is far more important for charity entrepreneurship than specific background experience or degrees.

Often, the number of diverse new skills you will have to learn will make any single background only slightly important relative to your ability to quickly absorb new things and develop meta-habits, such as task management. When individuals apply to our incubation camp, we try to inform those who could have been accepted if they had had more experience. And we aim to teach the individuals who get into our program the skills they would otherwise learn through many years of work experience. We will do it in a condensed and focused manner, so that more ground can be covered in a few months than in many years of more passive learning on the job.

Another option that we recommend is getting a **short-term job or internship**, should you want to gain skills during the year before applying for the CE incubation program. Taking on an internship between February and April can teach many skills and pass on many of the habits of an organization, yet still allow for a relatively quick founding of a charity after a minimal period of skill-building.

For the people who are positive that charity entrepreneurship is not a fit for them now, but think it might be in the future, the kind of work experience we most recommend is to work at an organization with roughly the following traits:

- Small nonprofit/charity
- Impact-focused

- Well-run
- In a more diverse role
- In the rough cause area you are interested in.

Small nonprofits will give you a better feel of working in CE directly (as your charity will start out very small), both work-wise and culturally. It will build a wide range of experiences, as you will often be required to wear a lot of different hats (a key skill for CE). Many small charities will end up giving you more responsibility within a shorter time frame. They also tend not to have the more established seniority-based systems that larger nonprofits often do, but which make it hard for an individual to grow. Culturally, they tend to be less rigid and have less established systems, leading to less rigid professionalism and more job ambiguity. Overall, working in smaller nonprofits will stand much closer to simulating the first few years of founding a charity than working in the larger ones.

Impact-focused charities will teach many of the skills that are hardest to learn through other work and volunteering experiences. Few charities are as single-mindedly focused on impact as the most cost-effective ones. An organization's focus will not only affect the skills you will learn (e.g. how to get website traffic versus measuring and evaluating your impact critically) but also the deeper attitudes you will end up building (e.g. it is easy to found a charity based around measuring the wrong metric).

Well-run charities will teach you habits that will benefit your start-up in the future, since you will often end up inadvertently replicating both the good and bad parts of the organizations you have worked for. This is part of the reason that **work experience is less important than people think**, since the habits you might pick up at a poorly run or non-impact focused organization can lead your charity down worse paths than coming in with a fresher slate for an incubation or founding setting. Well-run charities will allow you to pick up more cross-applicable meta habits, such as management, organizational structure, and task management.

Just as when working for a small charity, getting a more diverse role will teach a wider range of skills than getting a very specialized role. Managing a diverse team can also confer many of the same benefits. One exception to this is internships, even specific ones, that can teach a lot of specific skills in a short time span.

Generally, we emphasise that size and focus are more important than the specific cause, but it is a bonus if the charity is also working within the broader cause area

(e.g. animal issues, global health) in which you would like to found a charity in the future.

Some organizations that we think meet many of the above criteria include:

GiveWell-recommended or incubated charities:

- [Against Malaria Foundation](#)
- [Schistosomiasis Control Initiative](#)
- [Charity Science Health](#)
- [New Incentives](#)
- [Fortify Health](#)

ACE-recommended charities:

- [The Humane League](#)
- [Animal Equality](#)
- [The Good Food Institute](#)

Overall, we caution people against putting off their plans of founding a charity in order to build experience, but instead suggest they try to build experience in more rapid ways, such as through internships, incubation programs, and by talking to founders directly. That is, unless they have tried to found a charity and run into a block based specifically on needing more experience. As the next article discusses, many factors make now a particularly good time to start a charity – even with imperfect experience.

3. Should I Start a Charity Now or Later?

Author:

Ishaan Guptasarma,
Curriculum Developer & Content Writer

Is now the right time? Am I experienced enough? These are two major questions for many potential charity entrepreneurs. Here's our take.

When is the best time for me to start a charity?

We think that, for most people, the highest impact time to start a charity is as soon as possible. As a general rule, the earlier an intervention is implemented, the more good it does.

One of the most important benchmarks of impact is counterfactual speed-up, which refers to the degree to which you've caused something good to happen sooner than it otherwise would have, thereby positively impacting a greater number of lives. By starting a charity earlier, you help more people more quickly.

Additionally, some of the interventions we have researched are more time sensitive than others – they represent a chance to influence key decisions and steer specific events in a more positive direction. If we don't find the right people to implement these interventions, these windows of opportunity may expire.

It's also a uniquely good time to start charities in the context of working in the effective altruism community. The community has funding available, our incubation program is currently running strong, and there's a network of skilled alumni and mentors who are eager to help. These circumstances are all fairly recent developments and could change in the future.

The low number of people who passed the bar for being accepted into our program was the key bottleneck to our progress in 2019. Being able to accept one additional person would likely have done a lot of good for the world.

I'm not sure if I should start a charity because I feel inexperienced. Should I get more formal training, or pick up experience in other jobs?

Experience is helpful but not necessary. GiveWell, the Against Malaria Foundation, and many other high-impact organizations were sometimes started by people who had no formal background in the field. Work experience often does not fully prepare you to start a charity – dealing with operations for a large existing charity in the UK is very different from registering an organization and setting up a structure from scratch in India. Some people may come into starting a charity with twelve skills; others only start with eight. But if you need a hundred skills in total, that difference is far less important than your ability to learn new skills.

Our program includes people at all levels of education and experience. We train people with doctorates and people with over a decade of experience alongside those with little to no work experience or formal education. We have not yet seen that a co-founder's education and experience are strongly associated with success as an organization. When the need for specific skills arises, you can generally hire staff or consult experts who have training and experience in the necessary area.

If you agree with the fundamental principles of effective altruism, the most important element of starting a high-impact charity is choosing the right intervention. For example, an extremely competent organization that trains guide dogs to help the blind will likely not achieve as much impact per dollar as a moderately well-run organization that provides trachoma surgeries that prevent blindness (as the latter costs much less money). Finding the highest impact interventions requires systematically evaluating and comparing a very broad range of ideas. Subsequently implementing a high-impact intervention requires accurately updating your model of its effectiveness as new evidence comes in.

This means that the most important factor in starting a charity is not necessarily domain-specific training or experience, but rather the willingness to be flexible and follow the evidence where it leads – even if you have no personal incentive to do so. Not everybody has this quality. When you are running an organization, it is very easy to get caught up in other factors and lose focus on doing the most good, or to trick yourself into thinking that you are having a greater impact than you really are.

I'm not sure if it's the right time to start a charity because I have career plans, dependents, etc. right now. Can I still start a charity?

Probably! Many of the charity founders we know, including some who have succeeded within our program, have career plans, children, and so on. One of the founders we mentored founded Fortify Health a year before going to medical school, and the organization is still going strong. While starting a charity does involve a time commitment, it is possible to fit it in alongside other career and personal goals.

In general, we would suggest budgeting at least two years of full-time work to focus exclusively on starting your charity. If you think you are likely to leave the organization within two years of starting it, we advise having at least one co-founder who is able to commit to staying.

Charity entrepreneurship is a flexible career and rarely is there a single factor that rules it out as an option. That being said, being less flexible on some elements means you will likely have to be more flexible on others. If you need to work from a certain location, that will rule out some co-founders. If you are sold on a specific co-founder, they might only be a good fit for you with certain ideas. The more areas you are flexible in, the less important the other factors will be.

4. Plan B: Solid Alternatives to Charity Entrepreneurship

Author:

Patrick Stadler,
Director of Communications & Charity Mentor

Charity entrepreneurship is a high-risk, high-reward career option. Not all participants in the incubation program will become founders. And not all CE charities will succeed. It is important to acknowledge this from the beginning.

We at Charity Entrepreneurship consider starting a high-impact NGO to be one of the most impactful ways to make a change in the world. But we don't consider changing plans along the way to be a sign of failure. If you notice, for example, that your program is not having any impact based on a randomized controlled trial, it might be rational to give up at some point. Even during and after the incubation program, there are many good reasons to reconsider one's future as a charity founder. So it is smart to constantly reevaluate the entrepreneurship path and have one or multiple plan Bs ready.

Here are some of the reasons you might not become a charity entrepreneur:

1. There are many **drawbacks** to entrepreneurship, ranging from uncertainty to high risk of failure; from unglamorous work to a poor work-life balance. It is fair to recognize that you don't want to accept these drawbacks or that your personality might be too far away from the typical **traits** of an entrepreneur. Recognizing this early on spares you from a lot of hardship later.
2. Going through the incubation program you might not have found your ideal match in terms of a co-founder or charity idea. Given the importance of finding a **suitable partner in crime**, you may decide to hold off on starting an NGO.
3. Funding is a key bottleneck for charities. Some teams might not secure a grant after the program or struggle to get additional funds from third-parties. Such funding constraints can be overcome in the short term by working without pay or getting a paid job on the side. In the long term, however, you might not consider this sustainable.

4. As you work toward implementing a high-impact program in a particular country, you may notice that another organization is already implementing something very similar. Often that charity would not meet EA standards and does not have a focus on cost-effectiveness and monitoring & evaluation. Yet in some cases, you will come across a strong NGO already implementing your desired program or that could easily implement it based on its expertise. (For example, it might already have the infrastructure such as radio programming for health, so could easily add messaging for cause X). In such instances, you could first explore a collaboration with this organization or join their team rather than launching your own venture.

In sum, there are legitimate reasons not to pursue the path of entrepreneurship or to stop along the way.

Thankfully, there are fulfilling alternatives with a high impact profile. For some of these options, we would consider the expected impact of charity entrepreneurship to be higher. But that's on average and builds on very successful outcomes of a few charities (similar to the typical [power-law curve in venture capital](#) where a few companies make up a large part of a portfolio's success).

Here are some alternatives (in rough order of how close they are to charity entrepreneurship) and their unique advantages:

Working at CE

The organization Charity Entrepreneurship operates at a meta level, so staff members support not one but various charities. If you are a very research-oriented person, you will likely be happier at CE because it offers full-time researcher positions, while jobs at the average incubated charity lean more toward operations. Furthermore, at CE you are close to the incubated orgs and might be able to jump back into entrepreneurship at some point. These jobs are quite competitive but historically we have offered jobs to candidates who go through the program but have determined that entrepreneurship is not or is not yet a good fit for them. We also offer remote internships that can help test out an aspect of founding a charity in a lower commitment way.

- [Learn more about working at CE](#)

Working at a CE-incubated charity

Why found your own charity when you can join an existing NGO incubated by CE? Talent with an entrepreneurial mindset is in high demand at young organizations so you might be able to join a fast-growing organization as one of their first staff members, which gives you considerable influence. If you join an existing charity, the probability of success is likely higher as the initial battles are already behind it. You will still join a very dynamic organization but with more structure as some systems are already set up and co-founders will provide guidance. It is also one of the best ways to build up future skills for running your own project.

- [Learn more about CE's 2019 charities and Fortify Health](#)

Working at another EA charity

The advantages of working at an EA charity outside the CE network are very similar to those of working at one within it. The main difference might be that the charity is less influenced by the 'CE spirit' and more influenced by other streams within the EA movement. Look into:

- Animal charities such as [The Humane League](#), in roles recommended by [Animal Advocacy Careers](#) (AAC). Fundraising and operations roles are in [high demand](#). Additional details can be found in AAC's 2019 [talent bottleneck survey](#).
- Mental health charities, based on [career tips](#) by the Happier Lives Institute
- [GiveWell top charities](#) (or [GiveWell](#) itself)
- GiveWell-incubated charities such as [New Incentives](#)
- Research roles at impact-focused poverty research organizations e.g. [IDinsight](#), [J-PAL](#), or [3ie](#).

The importance of fundraising highlighted for animal advocacy applies to other areas too. We believe that in fundraising roles you can have particularly high leverage.

Working at another charity or foundation

Another option is to work for a charity in CE's recommended cause areas. To have the maximum impact working for these charities outside the EA movement, two factors should be looked for. First, that this organization is very well established and moves funding in the millions or tens of millions per year. Second, that your role is sufficiently influential to shape the organization's trajectory. These boxes are hard to check. Most NGOs are small, and at a large organization your impact

potential is often negligible. However, this can be an attractive option if you are able to steer a large organization in a positive direction.

Picking an NGO that already has some history of evidence-based programming might be a safer choice than battling windmills at one starting from zero. You could, for example, advocate for cash programming at an established development organization. Often these NGOs have some cash programs already but they make up only a small portion of overall programming. It's often easier to add a new program than to take one away. With many opportunities it will be best to add effective programs to the list of projects or supported organizations.

The same approach applies to **foundations**, where in the best case you could direct millions of dollars to evidence-based recipients. Foundations also have a hidden benefit of being able to donate a decent amount of money and thus double as an earning to give options.

Example: Why Working at a Foundation Might Top Earning to Give

- You make \$90,000 as a Program Manager at a foundation. You live off \$20,000 per year and turn your \$70,000 donation into \$100,000 through an employer donation matching program up to \$30,000.
- You direct a budget of \$1,000,000 that is on average 10% as effective as GiveWell top charities. With your inputs, the cost-effectiveness doubles (2x). You end up redirecting \$200,000 in terms of GiveWell equivalent funds. ($\$1,000,000 * 0.1 * 2$)
- Add the \$100,000 to the \$200,000 and you end up moving \$300,000 to GiveWell-like opportunities.
- This calculation improves further if you manage more than \$1 million, which is realistic at many foundations.

Working in policy

Several CE charities work on **policy**, yet the route of working in government is also feasible. You can pursue the role of a civil servant or politician/activist. This career generally provides more stability but requires persistence. One's impact should be calculated on the basis of decades rather than years. As national regulations impact millions and governments spend billions USD each year, your impact can nevertheless be considerable.

- [Check out this 80,000 hours article on working in policy](#)

Earning to give

While fundraising might be less of an issue in some areas (e.g. existential risk or animal welfare) there is still a demand for funding in general and particularly in areas such as global health. An earning to give (E2G) career can provide funding to jump-start many CE charities, e.g. through contributions to the Charity Entrepreneurship Fund or to Charity Entrepreneurship itself. While it's most effective if you focus on a [particularly lucrative career](#), you can still pick an area where you can maximize your strengths and job satisfaction.

As an extremely quantitative person, you might, for instance, enjoy working with complex models at a hedge fund more than getting into the weeds of start-up operations. So E2G might increase your job satisfaction compared to charity entrepreneurship.

Moreover, the period between the first funding and major follow-up funding is a difficult one for commercial and nonprofit start-ups alike. As an E2G donor, you can help charities move through this phase, commonly known as death valley.

E2G is generally a good option the more you can donate. If you can donate upwards of \$100,000 per year it becomes a very strong option. This is possible with a salary of \$135,000 or \$150,000 which applies to some jobs. Of course, the sweet spot would be a job where you make upwards of \$200,000 and could donate even more, which is a more limited set of jobs but can be a fit for some people. This has been achieved by some in the effective altruism finance community, for example.

Besides going for a high-salary job, you could also choose the path of business entrepreneurship, where your likelihood of success is lower but your earning potential higher. The entrepreneurship route might be one for those that realize that they love the founder's life but burn more for a for-profit topic, such as alternative meat. Many skill sets learned in this area can be applied to charity entrepreneurship and vice versa.

While [replaceability](#) is a concern for most EA concerns it is [less so](#) for E2G. There is no limit to positions that allow motivated individuals to donate large amounts of their salaries to EA causes (although there is a possible limit to the number of underfunded and cost-effective interventions).

Building career capital

Career capital is often best gained through a high-impact job such as charity entrepreneurship. There is a risk of **value drift** if you spend a few years in an environment that is not oriented toward alleviating the most suffering. In some cases, however, spending a few years in a larger, established organization might be helpful to jump-start a charity career. Strong management and operations skills with a focus on finances are, for instance, in high demand at EA organizations. Building skills and a network in fundraising can also come in handy, so that might be another area to gain career capital.

One advantage here is that you can focus on a narrower area of expertise compared to the broad field of entrepreneurship. This allows you to pick a field in which you excel and that you love (not necessarily the same thing). Depending on the industry, you might also be able to earn to give on the side.

Enjoying a higher **standard of living** is an advantage that applies to several alternative careers, specifically for E2G but also to career capital and working at foundations. These opportunities usually also have **higher prestige** from general society.

Summary

There are several suitable alternatives to charity entrepreneurship. Some of them might have a lower impact profile but suit you better. Some might also be a good interim solution until you jump back into your next adventure as a co-founder. The key point is that entrepreneurship is not necessarily a binary choice. You should not only consider the next few months, but also your overall career trajectory.

Key advantages of several alternatives to charity entrepreneurship

Alternative	Advantages
Work at CE	<ul style="list-style-type: none"> • Supporting various future charities • Research-only positions available • Option to jump back into a charity at some point
Work at a CE charity	<ul style="list-style-type: none"> • Higher likelihood of success • Dynamic but slightly more structured • CE spirit

Work at other EA organization	<ul style="list-style-type: none"> • Higher likelihood of success • Dynamic but slightly more structured
Work at other NGO or foundation in a chosen field	<ul style="list-style-type: none"> • Potentially much more established and high funding volume • Ability to shape non-EA org
Work in policy	<ul style="list-style-type: none"> • More stable environment (e.g. government servant) • Big lever • Higher living standard
Earning to give	<ul style="list-style-type: none"> • Focus on your key strengths and factors for job motivation • Funding gaps are prevalent for certain areas • Low risk of replaceability • Higher living standard
Career capital	<ul style="list-style-type: none"> • Get ready for future leadership role at an evidence-based charity • Some earning-to-give potential • Higher living standard

Exercise

Based on this article, how do you assess your plan Bs? Fill out a copy of the table below using the following rating:

- 1 → a top alternative: select a maximum of 1-3 alternatives here
- 2 → maybe, not your priority but still a feasible option
- 3 → not interested in this career track at this point in time

Explain your choice with a few keywords in a comments column.

Alternative	Rating	Comments
Work at CE		
Work at a CE charity		
Work at other EA organization		
Work at other NGO or foundation in a chosen field		
Work in policy		
Earning to give		
Career capital		

5. How to Increase Your Odds of Getting Accepted to the Charity Entrepreneurship Incubation Program

Author:

Kat Woods,
CE Co-founder

This handbook is part of a larger CE incubation program. Part of the course can be found online at our website. Part of the course is only accessible to founders who go through the formal [Charity Entrepreneurship Incubation Program](#). This program is however highly competitive, and many people reading this handbook might be interested in applying for it.

We're often asked what you can do to increase your odds of being accepted onto the Charity Entrepreneurship Incubation Program. While each person's answer will be different given their background and traits, here are the three most common things people can do (assuming you have already read this book):

1. Show that you can do things on your own
2. Teach yourself and practice good decision-making
3. Become an expert in effective altruism

Show that you can do things on your own

Most people's lives neither encourage nor support self-direction. Typical education models always tell you what to do, where to be, and how well you're doing. Same goes for the usual job, with a manager who will fire you if you don't do the things they tell you to do, to a certain standard, by a certain date. You may have some flexibility within that framework, but the scope for action is relatively narrow.

Entrepreneurship is entirely different. You are staring at a blank canvas. The only external accountability you have is in the distant future. You might only talk to a donor once a year. And you can't cram a whole year's worth of work into a week before you talk to them. It's not like school where you can get by with cramming if

you're talented enough. You need to do work every day even though nothing bad will happen in the immediate future if you don't.

What's more, there's nobody telling you which things need to be done in the first place. You could work on strategy, hiring, management, M&E, or even moving to Hawaii if you felt it was the best call. You have to pick the direction. Most people have little experience with autonomy. When they're faced with it, they're filled with immense discomfort at the uncertainty. That's why so many people postpone thinking about what to do after their education, often by simply getting another degree.

The good news is that these are all learnable. You just have atrophied initiative muscles due to disuse. All you have to do is practice. Once you do, the discomfort becomes smaller, and can be replaced by an exhilarating feeling of empowerment and freedom.

However, if you've never done it before, you may not be good at it. You have to learn how to motivate yourself when nobody else is helping you. You have to learn how to pick a good direction when there's no existing structure. That's why we look for people who have experience doing this. It's more likely that they'll be able to handle charity entrepreneurship: they've done this before and are not jumping into the deep end straight away.

Possible actions

- **Take an online course.** The simplest thing to do is start and finish an online course on something that you want to learn. This is relatively straightforward, but is a good way to dip your toes in the water. We give people a lot of points for listing an online course on their applications. Some courses that might be interesting and useful are:
 - **MIT course on program evaluation** by Nobel Prize-winning economists and founders of the randomista movement, Esther Duflo and Abhijit Banerjee
 - Abdul Latif Jameel Poverty Action Lab (J-PAL) **courses on evaluating social programs**
 - **Quantitative modelling**
- **Embark on an altruistic project.** Pick something that could be accomplished in one to six months that would make the world a better place. Start the project, and finish it. Ideally, use a good decision-making process to decide on what to do – the list in the section **below** will help you out. Not only will it

increase your odds of getting in, but will also reduce suffering in the meantime. Some examples might be running a community fundraiser for an effective charity or starting a campaign to convince your university to buy cage-free eggs.

Learn and practice good decision-making

Your success in life is determined by the direction you travel in and how efficiently you get there. However, often people focus on the latter, improving their capacity and productivity, while neglecting the former, thus getting nowhere fast. Making good decisions is a key factor in making sure you're picking the right way to go. This is crucial for charity entrepreneurship since, as mentioned above, you'll be facing a blank canvas in terms of what to do.

Many people are not very good at decision-making, their lives mostly characterized by bumbling around, stumbling upon things that are good enough. When asked why they chose a particular degree or career, they'll say, "I don't know. I guess I was good at it and liked it and I was accepted." Their process was opportunistic rather than deliberate.

Fortunately, decision-making is not a personality trait but a skill that can be developed. The broad two steps to follow for this are to:

1. Learn about good decision-making practices
2. Practice them

Learning is relatively straightforward. There are many resources on how to think about decisions. We've listed some below. Putting them into practice is harder. The biggest trick is remembering to do them in the first place. Unlike with some habits, usually there's no obvious trigger, since "make a decision" is hardly a concrete thing. Most of the time, making a decision doesn't feel like a decision. It just feels like life. However, there are a few situations where you can practise and hone your skills. These include choosing a:

- Cause area
- Way to use your spare hours at school or work to further that cause area (see the project above)
- Career
- Charity to donate to
- Personal budget

- Degree
- Job
- Skill to learn
- Place to live

Possible actions

1. Read the following:
 - a. [Using a Spreadsheet to Make Decisions](#).
 - b. How long to spend researching/learning compared to doing
 - i. [Optimal stopping](#)
 - ii. [Time capping](#)
 - c. [Steelman solitaire](#)
 - d. [Practice changing thought patterns](#). The post is about using thoughts to be happier, but can be cross-applied to any other habit of mind.
 - e. [Cluster approach](#)
 - f. [Algorithms to Live By](#)
 - g. [AI to Zombies](#)
2. Make friends with and surround yourself with people who are good decision makers. This is often the best way to internalize habits.
3. Apply these processes to the areas listed above.

Get a degree in effective altruism

We look for people who think well about how to maximize their impact using reason and science, which in essence means they are effective altruists. Much of this comes down to good decision-making, but a lot of it is also absorbing the lessons and thoughts that have already been discovered or expressed in the community. There's no need to reinvent the wheel.

While there's no way to get a degree in EA yet, you can still think of your knowledge level in EA as being akin to your level of education. Many people have an elementary knowledge of EA, having read a single book or watched a couple talks. Others have a PhD, where they've read practically everything there is on the subject and are on the cutting edge of a particular topic in the area. We're often looking for people who have at least a metaphorical undergraduate or masters. You can come in with a high school diploma if you have other qualities and skills that are sufficiently strong (and this goes for all criteria), but it will be a lot easier if you've got this qualification.

What does this look like? There are three different paths, and you'll ideally follow all three of them.

1. **Consume knowledge.** Read the various sources of knowledge in EA. Books, and forum and blog posts. Listen to podcasts, watch EAG presentations. Don't neglect to do the same for the particular cause area you're interested in. *There are plenty of resources outside of EA that are important to have for context.* There is no hard cut-off for what counts as a degree here, but around 300 hours is a good estimate.
2. **Act.** Put this knowledge into action. Book learning without practice is no use, as well as vice versa. Some things you can do are: volunteer, get an internship at an EA organisation, donate and practice thinking about where to donate, write your own material, start an altruistic project, lead projects at your local group, or start one if none exists in your area, or write an EA Forum post.
3. **Community.** Become part of the community. You will learn so much about EA from meeting fellow altruists. Being around them will passively increase your knowledge and actions through the topics of conversation and activities they participate in.

Possible actions

- Sources of knowledge:
 - *All of our articles*
 - *The CE reading list*
 - *EA Forum*
 - Try sorting by “most upvoted of all time” and working your way through
 - *EA Global conference talks*
 - *80,000 Hours podcast*
 - *Cause X Guide*
 - *Introduction – Peter Singer’s course on effective altruism*
 - Books
 - *The Life You Can Save*
 - *Doing Good Better*
 - *Poor Economics*
 - *Animal Liberation*
 - *Failing in the Field*
 - *Future Babble*
 - *Black Swan*
 - *Harry Potter and the Methods of Rationality*

- *How to Measure Anything*
- *Grit*
- Actions
 - Figure out where to donate and do so. Perhaps write up your thought process and share it with the community for feedback and inspiring others to do the same.
 - Start an altruistic project.
 - Volunteer for or get an internship at an EA organization.
 - Run events and projects for your local EA group. Run a fundraiser, host a presentation to build community (for example, *CE offers remote talks to EA groups*), or host an EA reading club.
- Community
 - Participate in your local EA group
 - Participate in the EA Forum and EA Facebook group discussions

In summary, there are many things you can do to increase your odds of being accepted into the Charity Entrepreneurship incubation program. Even if you don't get in on the first go, your life, skills, and impact will in all likelihood be improved by these actions. Hopefully this helps you, and we look forward to seeing your application!

IV. Co-founder

1. How to Successfully Pick a Co-founder

Author:

Patrick Stadler,
Director of Communications & Charity Mentor

No pressure, but picking your co-founder is one of the most important decisions you will ever make as a charity entrepreneur. Only selecting your intervention might be more crucial than deciding who will be your partner in crime.

This article addresses three key questions about your co-founder selection:

- Why is this decision so important?
- What should I look for?
- How should I look for it?

Why is this decision so important?

Each organization is, in essence, the people behind it. In the case of a start-up, that's the co-founders. Whom you pick as your partner will greatly affect decision-making and culture at your charity. Does your partner care about evidence? Are they energetic? These characteristics will reverberate across the organization as you recruit and manage staff, and interact with stakeholders.

There are many advantages to having a co-founder. Your partner will broaden your skill set, make your daily work fun, and increase the likelihood that you will stick through rough parts of the journey, as you can mutually support each other. Some start-up accelerators even have rules that recognize the comparative advantages of a partner: it is hard to get accepted at [Y Combinator](#) as a solo founder.

While you can reverse many decisions you make, the one about your co-founder is hard to turn around. A co-founder embodies part of the organization, so a departure can result in a loss of vision, implementing power, and atmosphere. It is no surprise that co-founder issues and split-ups are a **common failure mode** for start-ups in the for-profit world. In our experience, this finding translates to nonprofit start-ups as well.

What should I look for?

You got it, picking a co-founder is critical. So what aspects should you consider in your search process?

The ideal number of co-founders

Two is the sweet spot, according to many start-up connoisseurs. You get the benefits of a team, with less of the friction of larger groups. However, this is no scientific rule: the evidence in the literature is mixed and depends on the context. This **NYU study**, for example, concludes that solo founders outperform others in business settings. So these reflections should be taken with a grain of salt.

The friction of larger teams comes both in terms of decision-making and burn rate. The more people there are, the harder it becomes to make decisions, as there are more opinions and points of view. There is also the danger of coalitions forming within the founding team; for instance, two co-founders could habitually vote against the third member of the team. Rapid decision-making is a **key success factor** for a start-up, so anything that slows this down is a threat.

Besides decision-making, another factor that usually constrains charity start-ups is funding. Paying three instead of two salaries can be difficult at an early stage. The team might need to dedicate more resources to fundraising, rather than testing in the field. The pressure to scale an unproven concept will also be higher, as you might need to justify the bigger overhead to donors and to yourselves. With resources divided among more people, the salaries of each founding team member may also be lower than in the two-person scenario. This, in turn, could affect motivation or cause financial difficulties.

In sum, you likely don't want to be a solo founder, and you need some additional justification for having three or more co-founders as a charity entrepreneur. Such larger teams might be more common in the world of tech or pharma start-ups,

where very specific, hard-won expertise is required, and a team of three might combine business, tech and science backgrounds. In the world of nonprofits or common businesses, two members should be able to cover the basic skill-sets required. Historically and perhaps most relevantly, in our CE program two person founding teams have worked better than solo or three person teams.

Aligned goals and values

The brightest and most successful co-founder candidate cannot be a match if she or he does not align with your values and goals.

You both need to be fully convinced about the cause area and intervention that your charity pursues. Consider both the rational and emotional levels. An EA co-founder might be sympathetic to the idea of switching from a poverty to an animal charity from a rational perspective, but if he or she is not fully bought-in emotionally, it could get tricky in trying times. Similarly, you set yourself up for many arguments if your co-founder has never been fully convinced of the specific intervention area such as animal food fortification. Of course, the debate about the most effective intervention is vital at the beginning of a venture. But it should be settled within an agreed timeframe, and should not constantly pop up during daily operations. You will hit many roadblocks with any intervention, and this should not be an excuse to switch back and forth.

You need to share the same personal goals and timeframes. It will be a recipe for conflict if you see yourself leading a charity for 5-10 years but your co-founder is interested in “trying out” charity entrepreneurship for two years. Watch out for co-founders who don’t understand that setting up a new organization takes years, or who are already planning their next big thing.

You should share the key values of evidence-based charity entrepreneurship and effective altruism. Cost-effectiveness, solid impact measurement, tractability, and transparency all fall under this category. Don’t trade these values for expert status or experience. Such compromises can be made at the level of your field staff, but at the co-founder level, ignoring these values will lead to poor decisions or **value drift**.

Grit: a short word with profound implications. Pick a co-founder who has shown grit and tenacity in his professional and private life. Start-ups are often uncomfortable and hard, and you don’t want your co-founder to jump ship when the first clouds appear.

There are many intelligent people, but your co-founder needs to combine intelligence with a proactive attitude. Pick someone who can get sh... done over the smartest person in the room. The journey of a start-up involves a lot of trial and error, and only moving forward will help you get the necessary feedback on your journey to success.

As in relationships, financial preferences are a common cause of conflict between co-founders. Ensure that your co-founder has similar ideas on expenses. Lean toward frugality, as a high burn rate limits your time to test your program, while a modest budget attracts donors.

Pick someone who shares a similar perspective on work-life balance. There does not need to be perfect overlap here, but avoid going for someone who prefers to work much more or less than you, as this will lead to frustration or burnout. Of course, you can adapt your idea of work-life balance as your organization matures, but it helps to have a shared general understanding at the beginning.

Complementary skills

You both graduated from an Ivy League school in business administration, both like to work on headquarters level tasks, and neither of you has ever worked in a developing country. That might not be the best setup in terms of complementary skills and interests. Ideally, one co-founder brings experience and interest in working in the field, in the case of a poverty charity. It is also helpful to have previous involvement in the area you are working on, for instance, it is beneficial if at least one co-founder at an animal charity has been engaged in the space. It is common in for-profits to have a technical founder who is responsible for software and product, and a non-technical founder responsible for the business side. At nonprofits, a well-balanced team can often consist of a founder with a background in research/programming and another with a background in business/communications.

Certainly, it's unlikely that you will achieve perfect overlap in terms of complementary skills. But strive to build a team where you have unique strengths and weaknesses. In the best case, you understand who is going to lead which domain from the beginning. Someone might be uniquely qualified for fundraising and other HQ-level tasks, while the other excels at program implementation in the field.

Consider the following skills and interests as you look for complementarities:

- Cause-area experience
- Cost-effectiveness analysis
- Budgeting and financial planning
- Operations
- Fundraising
- Stakeholder relations
- Communications
- Human resources (incl. recruitment)
- Managing people
- Planning and task management
- Monitoring and evaluation
- Research
- Negotiations
- Customer focus and user experience

Complementary psychology

Potentially more important than complementary skills are psychological traits that align within the co-founding team. This does not mean that you should be twins – it might well mean that you have opposite psychological traits. Someone might be more easy-going, while the other always has strong opinions. Someone might be more enthusiastic about new ideas, the other more skeptical and suited for reality checks. Someone might have many ups and downs, while the other is steadier. The team goes beyond its members.

On the other hand, differences in psychology can cause conflict. If one partner gains energy from heated disputes while the other would prefer to discuss disagreement more calmly, you need to find a way of dealing with these differences.

All sounds a bit complex and fuzzy? Just consider whether you enjoy spending time with, and tackling challenges with, the person in front of you. If so, you might have found a good match from a personality perspective. As a step further, imagine that your charity has failed after multiple years of hard work. Would you still have enjoyed the journey together as co-founders? There is no need to be best friends, but don't go into this if you need to justify the decision to yourself ("we might not click, but the opportunity is too big to pass up..."). That's a bad start, as in the end, a venture is all about people.

Logistical factors

So, you have found your perfect match but he or she lives on a different continent. Is this a dealbreaker? Start-up wisdom holds that you should be based in the same location. In-person communication often allows you to move faster and build stronger personal ties. Yet, there are examples of business (e.g. Zapier) and nonprofit start-ups (e.g. New Incentives) that had a remote setup from the beginning without experiencing substantial disadvantages.

If you are looking for a co-founder, you want to still prioritize someone with whom you can work in person. This includes candidates who do not already live in your location but are willing to relocate – and, importantly, who have the necessary citizenship to get a visa. Unfortunately, this last factor restricts your choice and is unfair to potential founders from certain countries. Yet, experience shows that visa status should be considered. Several previous Charity Entrepreneurship graduates struggled with this issue.

Real work

As with recruiting staff, the impression of a candidate from their CV and conversations might be misleading. The real test only comes with real work. As a result, you would want to test your co-founder working relationship with minor common assignments before you jump right in. You could, for instance, work on a short research summary or an overview of the planned intervention. It's a great sign if such a test collaboration is inspiring to both sides and results in a strong outcome that goes beyond what you each would have achieved as individuals.

The longer you have known a person, the more certainty you have about their suitability. However, if you're considering starting your charity with a good friend, consider the following questions:

- Has your relationship included common work assignments, or was it limited to fun leisure activities? Of course, you would want to consider how the other person performs in a work situation.
- Are you fine with your start-up impacting your friendship? If you launch a venture with a good friend, both of you should acknowledge that this can go incredibly well, as you know and respect each other, but it can also mean the end of a dear friendship. Be sure you both understand that in case of conflict,

the charity's future, not your friendship's trajectory, will likely be the deciding factor.

How to look for your co-founder

You know about the importance of a co-founder and their ideal traits. Now you just have to find your match. Easier said than done. Yet there is an effective and comfortable way to find a co-founder: join [Charity Entrepreneurship's two-month Incubation Program](#). The incubator brings together talented individuals, gives them training in all aspects of running a charity start-up, and funds their high-impact ventures.

The program prioritizes various activities aimed at successful co-founder selection. It starts with fun team-building activities, ranging from solving moral dilemmas to taking personality tests to recounting each other's life story. As a next step, participants work with different partners on actual start-up tasks, such as drafting a fundraising plan or developing a monitoring and evaluation (M&E) concept. What's more, you will get to know potential co-founders in after-work activities during the two-month program. So you will get a strong understanding of your co-founder before you make the decision that defines your charity, and potentially your professional life.

If you plan to start your charity outside the Charity Entrepreneurship program, we recommend modeling your process similarly.

- Make meeting potential co-founders a priority, and go about it in a systematic way. Don't expect the first person you meet at a random event to be your match made in heaven.
- Go through the list of characteristics above together.
- Make use of questionnaires aimed at picking a co-founder. FirstRound has published such a [questionnaire](#) – it is geared toward for-profits, and includes some questions that don't apply to a charity but it is still a helpful resource.
- Respond to more general questionnaires, such as these [36 questions](#) that originate in the field of dating but add value when picking a co-founder as well.
- Take personality tests to identify potential areas of overlap, complementarity or friction. [Openpsychometrics](#) is a great resource in this regard.

- If you are confident that you found a match, start discussing a Founders' Agreement, as outlined in this [article](#) on strengthening your co-founder relationship.

Having a clear process and defining what you are looking for is a solid starting ground. At the same time, be realistic. You are not looking for a superwoman or superman, but a human being who has flaws like you. Ask about whether you complement each other well, and whether both or all potential co-founders are a good fit for (charity) entrepreneurship. So instead of finding someone who stands out in all areas, define your own set of co-founder must-haves and nice-to-haves as you go about your search.

If you don't immediately find a co-founder you are looking forward to working with, give it some time. You can always start working as a solo founder, and bring someone on board later. This still trumps settling for a poor choice. Eventually, you will find your perfect match.

Happy searching!

Summary

- **Why is it important?** A co-founder complements your skills, makes your work more fun and helps you stick through hard times. The choice is important as you cannot undo it easily and co-founder friction is a common cause of start-up failure.
- **What do I look for?** In most cases, aim for a team of two co-founders. Your partner in crime should have: aligned goals and values (e.g. cost-effectiveness, proactive attitude), complementary skills (e.g. the field experience you lack), complementary psychology (e.g. enthusiastic if you are skeptical; similar conflict resolution behaviors), and a visa to work in person as a co-founder team. Common work experience is helpful too as it gives you a real-life sense of how well you jive, so start with some trial tasks if you have never collaborated prior to your start-up.
- **How do I look for it?** The two-month [Charity Entrepreneurship Incubation Program](#) is an ideal setting. It includes various team-building exercises and real-work projects that help you find your ideal co-founder before launching your charity with mentorship and funding. If you plan to start your charity outside the program, design a similarly systematic co-founder search process. Go through the criteria in this article, take the co-founder

questionnaire, compare personality tests, and discuss a Founders' Agreement.

Reading list

- FirstRound's [Founder Dating Playbook](#)
- Excerpt from Sam Altman's Startup Playbook: [A Great Team](#)
- [Empirical data on value drift](#)

2. Working with Your Co-founder: How to Excel at Joint Decision-making, Task Management, and Communications

Author:

Patrick Stadler,
Director of Communications & Charity Mentor

Now, work can start. You picked a co-founder with shared values and goals, a complementary skill set and compatible psychology. Ahead of you are weeks, months, and years of work in scaling your charity from a small start-up to an established organization. How you work with your co-founder will be decisive. This article outlines basic lessons for successful collaboration in day-to-day work, while the following article sheds light on how to strengthen the relationship with your co-founder at a deeper level ([How to strengthen your co-founder relationship](#)).

Communications

Successful communication is at the core of every task you tackle and decision you take. As outlined in the previous article, working in the same location is beneficial, especially as your organization starts out. Understanding the communication preferences of your co-founder is helpful too. One might be chatty while the other mostly prefers quiet time. Find arrangements that suit both styles, for instance, by defining deep work hours without interruption, semi-deep work hours, and normal work hours during which interactions are fine. The rules for these three distinct work modes could capture both personal communication and instant messaging (e.g. snoozed notifications during deep work). Consider that original work needs uninterrupted space on the calendar, as Paul Graham outlines in his classic essay [Maker's Schedule, Manager's Schedule](#).

Communications, of course, goes beyond such logistical factors. At its deepest level, communications affect your psychology as individuals and thereby influences your decision-making. A look at communications theory and research on couple's therapy can be illuminating.

Dr. John Gottman is famous for predicting the longevity of romantic relationships by assessing videos of interpersonal arguments. He identifies **Four Horsemen**, inspired by the biblical image of the four horsemen of the Apocalypse, that indicate the deterioration of a relationship. They include:

- *Generalized criticism* (“You never...”)
- *Defensiveness* (“I don’t have any responsibility...”)
- *Stonewalling* (“I am not willing to talk about this...”)
- *Contempt* (“You are unable to contribute anything meaningful...”)

One key idea behind Gottman’s research is that negativity has a profound impact on a relationship. He thinks of it as an *Emotional Bank Account* to which every positive or negative action adds or subtracts. The challenge is that every negative action weighs more than several positive ones. Besides avoiding negative interactions, you, therefore, want to proactively work on adding positive interactions. This can include showing more appreciation with sincere compliments, which is especially important as we often take things for granted in daily life. Don’t wait for the eulogy to commend your co-founder! Planning fun stuff outside your work relationship can help **strengthen your relationship** too.

The theory of non-violent communications developed by Dr. Marshall Rosenberg in the 1960s offers a few valuable principles on staying away from the Four Horsemen. The core idea is that we all have universal human needs that we can express in a positive way unless we feel threatened. This two-minute **video** provides an overview of how the theory applies in practice. Four elements make up a statement in a non-violent conversation.

- *Observation*: a factual observation specific to a particular time and context; without any evaluation
- *Feelings*: showing vulnerability by expressing your feelings connected to your needs (real emotions not thoughts)
- *Needs*: your underlying need that is affected
- *Request*: request for a specific action in a clear and positive language (not a demand)

The usual structure is “When (observation) then I feel (feelings) as I have a need for (needs) so I would appreciate it if (request)”. Or in a specific example: “When you are late I feel neglected because I want to use my time well. I would appreciate it if you could let me know when you are running late.”

As often in life, this is a simple concept that is not always easy to apply. But you don’t need to become an overnight master of this concept, and you will also likely

never be able to adopt it in 100% of your social interactions. A gradual approach with increasing adoption by you and your co-founder can, however, go a long way.

Task management

You don't need the latest sophisticated task management methodology or tool to succeed. At the same time, translating your vision into actions requires planning and prioritization.

Here are some minimum guidelines that you and your co-founder should follow:

- Keep it simple: unless you are both really into productivity hacks, a simple task management method will do it. The key is to stick to it.
- Works for both: this might be obvious but of course, whatever system you adopt has to work for both of you.
- Short and long term: ensure that your task management considers both short- and long-term tasks. Longer-term goals often get neglected in the daily hectic.
- Important vs. urgent: a solid, simple framework is to plot out your tasks into four quadrants by importance and urgency. Focus on the important tasks, starting with the urgent, important tasks. Avoid the temptation to spend too much time on the urgent, unimportant tasks. These are the typical “operational fires” that come up on a daily basis: for instance, a field staff member not being able to work without access to an IT system. Establish standard processes to delegate such tasks to other staff or freelancers. Also stay away from the non-urgent, unimportant little tasks. This is sometimes tricky as solving them can give you instant gratification, e.g. looking into a redesign of your logo which is neither urgent nor important.

Getting Things Done (GTD) is a popular method for time and task management. It might also work for you. Some good places to start are this [15-minute guide](#) on GTD and this [minimalist time-management system](#) by the Operations Director of the EA Foundation.

In terms of specific tools for task management, it's also a matter of taste. Some might go for a simple Google Sheet. This could include a GANTT-style template for planning over months and a list for short-term tasks. The advantage of this approach is that it integrates well with your other sheets and documents that you likely host within Google Drive. The downside is that it offers less task management specific features such as reminders. Asana and Trello are two popular options for

task management software. Look out for nonprofit discounts. As an EA charity, you might get a special discount for Asana through the EA Hub.

Decision-making

Prioritizing your tasks and time requires a multitude of decisions every day. In fact, one could rightly say your main job as co-founders is to take decisions (which is why [part two](#) of this handbook covers this topic in such depth). How to go about this, especially if you have different intuitions about the right path forward?

First of all, be aware of common issues in decision-making and basic methods of arriving at better decisions. As Nobel prize winner Daniel Kahnemann has shown in his work, we have consistent biases in decision-making. We are, for instance, generally loss-averse and therefore value what we could lose higher than what we could gain. The loss of \$5 feels worse than the gain of \$5. This [cognitive bias cheat sheet](#) provides a great visual overview. Good to consult occasionally to remind yourself.

Kahnemann also points out that [noise might even be worse](#) for decision-making than biases. While biases are predictable and uniform, noise is random and therefore harder to spot. As one study suggests, radiologists give the same x-ray a different diagnosis in 20% of cases. This reflects the randomness of noise rather than the uniformity of a bias. One way to counter this is to set up strict processes, ideally algorithms, for your most important decisions. Sounds complex? A good start is to use spreadsheets as a powerful decision-making tool, for instance, by listing and weighing your most important criteria before you act.

Second, it helps if you have outlined a general framework for your decision-making. If you approach a decision with the same values such as cost-effectiveness you are off to a good start. You are already thinking along the same lines. It also helps if you have outlined some roles and responsibilities. Many decisions don't necessarily need the approval of all co-founders, which makes for more rapid decision-making and clearer accountability.

Third, if you disagree despite all of this, you still have a variety of options at your disposal. Here are a few specific ways to resolve deadlock around a decision:

- What does the data say? Disagreement is usually around different intuitions to interpret a specific challenge. This can be resolved if you have data pointing in one direction or the other.

- Can you test? Even if you don't have existing data to resolve the controversial question, you can potentially set up a short experiment to get answers. If you target fish farmers as an animal charity and disagree internally about the best messaging, test two or more different versions. You might not always be able to run a statistically significant test but your experiment can at least give you a better understanding from a qualitative research perspective.
- Can you reach a compromise? It does not have to be A or B, maybe the concerns of both co-founders can be integrated into a solution AB.
- What are the underlying concerns of the dispute? Dive one level deeper and try to understand the concerns of each co-founder. As outlined in the classic book on negotiation strategy *Getting to Yes*, focus on the interests, not the positions. You might have a disagreement, for example, about a lower/higher reimbursement rate of field staff. The person who insists on the upper bound might not necessarily be stuck to a position of higher reimbursement rate. Her underlying interest might be to adequately value the contributions of the field team. This could also be solved by other means, e.g. social recognition or a performance-based bonus. Focusing on the interests opens up the field for various out-of-the-box solutions.
- Can an expert settle the dispute? You might both agree that your advisor X or your advisory board is best suited to solve a particular conflict. In this sense, you agree to the decision-making method of delegating to someone with more expertise or experience.
- Can you defer the decision? While rapid decision-making is a key advantage for start-ups, not every decision needs to be taken right at this moment. It might be possible and advisable to wait for a few more days, weeks, or months. Sometimes a decision becomes irrelevant in the meantime or you both converge to one solution.
- Can you include the legal board? This should usually be a last resort, reserved only for strategic high-level decisions or a case of severe co-founder conflict. Legal boards only meet a few times per year so they are not the best forum in which to take most decisions. Best to reserve the legal board to weigh in on strategic high-level decisions all co-founders agree on or standard oversight tasks such as auditing.

Communications, task management, and decision-making are key components of working with your co-founder. If you would like to strengthen your relationship further, you should also consider regular happiness and collaboration check-ins and draft a Founders' Agreement, which we turn to in the next article.

Summary

There are three essential components of working with your co-founder:

- **Communications:** build your communications culture around your preferences. Factor in sufficient time for concentrated work (maker's schedule). Avoid negative patterns in your communications, as outlined in the Four Horsemen by Dr. Gottman. Instead, implement the four principles of Non-Violent Communications as developed by Dr. Rosenberg.
- **Task Management:** a simple, consistent time and task management setup trumps a sophisticated one you never implement. Organize your tasks around importance and urgency. Stay away from instant gratification tasks that are often unimportant and not urgent, and delegate unimportant urgent tasks to your staff. In terms of tools, consider Google Sheets or a dedicated task management app such as Asana or Trello.
- **Decision-making:** implement general guidelines for good decision-making such as using spreadsheets for listing/weighing criteria, avoiding biases, and spotting noise. Shared values and some basic understanding of your roles and responsibilities provide a solid ground to make decisions. If nevertheless you disagree, run a test, compromise, find consensus by looking at interests instead of positions, or defer to an expert.

Reading list

- [Maker's Schedule, Manager's Schedule](#) (Paul Graham)
- [Beginners Guide to Gottman Method](#) (Four Horsemen)
- [Video: What is Non-Violent Communication](#)
- [Bias and Noise: Daniel Kahneman on Errors in Decision-Making](#)
- [How to work together](#) (Y Combinator Startup School)
- [Getting to Yes by Robert Fisher, Summary](#)

3. How to Strengthen Your Co-founder Relationship

Author:

Patrick Stadler,
Director of Communications & Charity Mentor

“Your co-founder relationship is like a marriage.” This common statement from experienced start-up operators might sound like a stretch, yet there are undeniable similarities in these relationships. Like spouses, co-founders spend considerable time with each other, but instead of taking care of children, you are nurturing your upcoming organization which, in the case of charity entrepreneurship, could have a profound impact on the world.

Your co-founder relationship might well be one of the best things in your life and be a source of energy for years. Unfortunately, as in the case of marriages, co-founder conflict or split-ups are more common than people assume. In fact, they are one of the most common reasons that start-ups fail.³ Need to pivot and change your intervention? Need to work without salary for a few months? That’s all doable if the team remains committed. However, if co-founders don’t get along, that often is the beginning of the end.

How do you ensure that co-founder issues don’t come up in the first place? This might not always come naturally. As a start-up team, you are enmeshed in the daily grind of testing and delivering your intervention. Anything longer-term, such as relationship building, gets forgotten in the most challenging periods if you don’t make it a priority.

Thankfully, there are a few ways that you can foster your co-founder relationship without burdening your operational priorities. The following have been especially helpful:

1. Founder nights
2. Happiness & collaboration check-ins

³ Linkilaw (2019) ‘5 Reasons Why Co-founders Fall Out’ Available at <https://linkilaw.com/start-up-advice-tips/5-reasons-why-co-founders-fall-out/> (Accessed 10/30/2019)

3. Founders' Agreement

Founder nights

All work and no play? Planning for some fun time together can make sense, especially during phases of stress and pressure. This ensures that your interactions are not limited to solving operational issues and making tough decisions. Depending on your interests, you could go for a simple dinner, or to the movies, sports activities, or a weekend getaway. Avoid the temptation to discuss running your start-up. Focus on being in the moment, sharing stuff from your private life, and discussing broader topics that excite you.

While the key benefit here is bonding, you will notice that these casual meetups can also provide some interesting insights. You might, for instance, discuss the latest trends in technology or findings in psychology and notice that a particular approach could also be applied at your organization. But again, the key benefit is relationship building, so just sit back and enjoy.

As artificial as it might feel, make sure to schedule and prioritize these bonding nights. Otherwise, they will often get forgotten, or canceled due to other pressures.

Happiness & collaboration check-ins

Going out for a movie once in a while is not sufficient for a lasting co-founder relationship. While the founder nights are purely about enjoying companionship, the happiness & collaboration check-ins allow you to discuss your satisfaction levels as a co-founder pair and as individuals.

One option is to include these topics in a regular 1:1 meeting that you have scheduled anyhow. However, it might be best to clearly distinguish this type of meeting and schedule it regularly. This prevents you from skipping over the happiness & collaboration questions due to urgent discussion points. Moreover, it sets a different tone and helps you separate this discussion clearly from tough conversations about daily implementation or strategic priorities. Such a dedicated happiness & collaboration check-in could, for instance, take place on a monthly basis for two hours.

These are some of the questions that can guide your happiness & collaboration check-in. As you can see, the list starts with individual satisfaction, including private life, which also affects job satisfaction. It then moves to the interpersonal

level, the co-founder collaboration. The questions intentionally include references to common successes, vision, and gratitude. Those are otherwise often neglected, as we are used to focusing on areas for improvement. The separation between quick wins and long-term solutions allows for immediate progress, acknowledging that some differences between founders might not be easily resolvable:

- What am I grateful for personally and in our co-founder relationship?
- What has been my personal satisfaction level outside of the job recently?
- What has been my personal satisfaction level at the job recently?
- What common successes did we achieve recently?
- What are some aspects that could be improved quickly (low-hanging fruit)?
- What are some aspects that could be improved over time?
- What is our vision?

You can decide whether and how to track the progress of these conversations. In general, capturing the outcome of meetings in a Google Doc is a good idea. Yet with these check-ins, it might feel more natural to keep it to a discussion where only a few or no notes are taken. This might also help you open up as a team and tackle the underlying issues of potential conflict.

Founders' Agreement

Intense and legalistic. These are two common typical reactions when aspiring entrepreneurs first hear about the idea of a Founders' Agreement (FA). Such a FA outlines the expectations and preferences of co-founders on the full spectrum of their working relationship and project.

The idea here is not to draft a static contract that manages every aspect of the relationship in detail. An FA should rather be seen as a dynamic document that triggers the right conversations between founders and prevents misunderstandings and bad feelings. As such, it goes a step further than the happiness & collaboration check-ins and helps manage the relationship on a more strategic level.

What goes into the FA? This depends heavily on the co-founder pair. It usually involves key questions on your collaboration: questions that are often not openly discussed, (wrongly) implicitly assumed, and in some cases divisive.

- **Roles and responsibilities:** Your roles and responsibilities might be in flux, especially at an early-stage start-up. However, the more you can define them, the better, as it will help create ownership and prevent misunderstandings (“Oh, I thought you were responsible for this!”). You can

define roles and responsibilities top-down or bottom-up. Top-down would mean starting with core functions of your organization such as fundraising, program development, and recruitment. Bottom-up would be looking at 10 specific tasks that recently came up, and seeing whether you agree on who is responsible.

- **Decision-making:** Related to roles and responsibilities are your decision-making procedures. Initially, it makes sense to take all major decisions on a consensus basis. As your responsibilities get defined more clearly and the organization grows, you can hand more of the decision-making to the related co-founder. Strategic decisions can continue to be taken as a team. Some of the points you could address in an FA: which decisions do we take together, which do we delegate to a co-founder? How do we deal with lacking consensus? (e.g. include third-party, let an experiment determine the course of action, ...)
- **Field time:** Being close to your customers is key in an early-stage organization, as it provides you with critical feedback and allows you to test different iterations of your intervention. For many organizations, this involves spending time in developing countries, which can be draining and challenging for founders who might be separated from their families and friends. From an organizational perspective, you want to optimize for as much field time of your leadership team as possible. From a personal happiness perspective, however, there are limits and those should be considered. It is important to make those preferences explicit and find a solution that considers the field presence needs of the organization as well as personal preferences which indirectly affect the success of the organization. Specific agreements might include: how often do you spend time in the country? Where exactly are you based (capital vs. regional office)? For how long do you stay at a time?
- **Work/life balance:** Do you work seven days a week for 12 hours a day? Or do you want to enforce email-free weekends? As a start-up, you are likely going to invest more than the regular 40 hours a week. Yet personal expectations might differ and should be discussed openly. There is a clear case for not overworking yourselves, as this can result in burnout and negatively affect your decision-making. Specific points to address: do you track co-founder's work time and if so how? Are there any work limits? Are weekends intentionally work free? What is your co-founder vacation policy?
- **Communications:** Strong communication between co-founders is key to your success. In the FA, you may also address topics such as communication preferences (time, channels, difference per topic). You might, for instance, define that non-urgent topics are generally discussed in Slack while urgent

questions are sent on WhatsApp. Or you might want to generally limit the distracting flow of instant messages and emails and write down all non-urgent questions for your co-founder on your next meeting's agenda. This is especially helpful if you try to adhere to a **maker's schedule** that avoids distractions. Such rules might seem trivial but can improve work happiness and efficiency considerably.

- **Personal happiness and growth:** A team is only as good as its members. Define how each co-founder is committed to taking care of her/himself through sports, meditation or whatever works for her/him. Outline areas for the personal growth of the founders explicitly. This ensures that the founders grow as the organization grows and longer term development does not fall victim to daily operational tasks. Founder A might, for example, want to further develop his Google Sheets and statistics skills while founder B might focus on leadership development and public speaking classes. Growth areas could build on strengths or tackle weaknesses; both strategies have their pros and cons. A lot will depend on the organization's trajectory and personal priorities.

As you can see, an FA is flexible and builds on the preferences of you, the co-founders. Make a copy of our **template** to get started, and feel free to add or leave out topics. And remember that the FA, while giving some stability and consistency to the co-founder relationship, is not set in stone. Revisit the FA every 3, 6 or 12 months to see where you stand and adapt it accordingly.

A final point: an FA is a great foundation to draft your **Manager User Guide**. Such a user guide defines your preferences as a manager and helps your broader team understand and work with you better. A FirstRound **article** sheds light on the steps to arrive at your own Manager User Guide.

Founder nights, happiness & collaboration check-ins and an FA – these are three methods to foster your co-founder relationship, the nucleus of every successful charity start-up. Tailor them to your personal preferences, be sure to implement at least some of them, and the likelihood of a strong and lasting co-founder relationship will increase.

And remember: a co-founder relationship is often one of the most profound and satisfying relationships in your life, so enjoy and appreciate being on a mission together!

Summary

- Many start-ups fail or don't execute at their highest levels due to conflict among founders. Yet strengthening a co-founder relationship often gets neglected amid the daily grind, especially during difficult times.
- These three methods help you to sustain a healthy co-founder relationship. While this is not rocket science, you need to make it a priority for it to actually happen:
 - **Founder nights:** hanging out as co-founders without discussing work
 - **Happiness & collaboration check-ins:** 1-on-1 meetings specifically to address personal/group happiness and collaboration and related room for improvement
 - **Founders' agreement:** a written framework reflecting your preferences as co-founders that guides your collaboration and gets regularly updated. This is not a static or legalistic document, but rather a tool to help you raise important questions and provide clarity about the guidelines of your partnership.

V. Community

1. Introduction to Effective Altruism

Author:

Ishaan Guptasarma,
Curriculum Developer & Content Writer

Charity Entrepreneurship is part of the effective altruism movement. If you're going through our incubation program, it's probable that some of your potential donors, employees, partner organizations, and other stakeholders come from the effective altruism community. The first article in this chapter orients you to effective altruism; we'll then turn to the community formed around CE itself, and wrap up with some articles that look at cause-specific communities.

While the basic premise of effective altruism seems simple, the set of ideas considered "common knowledge" among people heavily involved in effective altruism can be quite complicated. In some ways, we might conceptualize effective altruism as a field of study – it has a few frequently cited influential thinkers, a set of widely read canonical works, an internal jargon, and several diverging schools of thought. If you're new to effective altruism, becoming well-versed in this contextual knowledge will help you communicate with other people in the movement more easily. The movement is full of incredibly talented people many of whom are in good positions to help early-stage organizations.

EA Hub's [About Effective Altruism](#) page has a collection of videos, books, blog posts, podcasts, and online courses geared toward explaining the basic principles of effective altruism and persuading people of them. If you're very unfamiliar with effective altruism, it's a good place to dive in.

Since you're creating an organization, it's especially important to be aware of [organizations associated with effective altruism](#).

If you understand the basics, you can keep up with current events by subscribing to some [newsletters](#), [podcasts](#), and [Facebook groups](#). Wherever you're reading this from, you should also consider connecting with a [local group](#) if there is one in your area. The [EA Forum](#) is a good place to engage in active dialogue with the community, and to post ideas for feedback.

Our [chapter on epistemology](#) aims to bring you up to speed on some of the more complex ideas and the key ideas of EA that are most important to apply to founding new charities. However, it will broadly assume some knowledge of the community.

Another valuable community worth knowing a little bit about is the [rationalist movement](#), which overlaps somewhat with the effective altruism movement. Most of the core concepts can be found by reading [The Sequences](#). We will also link key articles in our classes when they are uniquely relevant.

If you've familiarized yourself with the ideas, culture, community, and major institutions within effective altruism and are looking to get more advanced, it can be useful to learn the history of the movement and relationships between organizations and how some of the underlying [funding structures](#) work. Some useful things to learn about include the history and relationships between [GiveWell](#), [Good Ventures](#), [Open Philanthropy](#), [Giving What We Can](#), [Center for Effective Altruism](#) (including [Effective Altruism Funds](#)), [80,000 hours](#), [Rethink Charity](#), and Charity Entrepreneurship. Charity Entrepreneurship, like many other organizations affiliated with effective altruism, receives grants from [Open Philanthropy](#) and [EA Funds](#) as well as from individual donors.

Resources

Introductory

[EA Hub – Learn about Effective Altruism](#)

Intermediate

[EA Forum](#)

[Sequences](#)

[EA Survey 2018 Series: Community Demographics & Characteristics](#)

[EA Survey 2019 Series: Community Demographics & Characteristics](#)

[The EA Community and Human Interventions](#)

The EA Community and Animals

Advanced: History and relationships of key organizations and funders

 Major Funders within Effective Altruism
Separating GiveWell and the Open Philanthropy Project
Good Ventures: Our Portfolio
Open Philanthropy: Who we are
Giving What We Can: Our History
80,000 hours: About us.
Center for Effective Altruism: History
About Effective Altruism Funds

2. The EA Community and Animals

Author:

Ishaan Guptasarma,
Curriculum Developer & Content Writer

If you plan to create an organization that helps animals and wish to recruit the support of the effective altruism community, it helps to have a good sense of what the rest of the community is thinking and talking about. We've compiled a series of articles to help you keep up-to-date on the conversation.

The articles that we most recommend you read are marked in **bold**. Other articles may not be necessary to read unless you are particularly interested, or working in a relevant area.

Overview

Effective Animal Advocacy Resources by Rethink Priorities: This is a fairly comprehensive list of animal resources in the effective altruism movement. If you are working with animals, we recommend being somewhat familiar with all of the organizations and fundraising bodies listed here. It also includes various newsletters, repositories, notable articles, and notable people, which may be interesting. A great place to start orienting yourself to the community. If you're trying to get in touch, there is also the **Effective animal advocacy community directory**, a partial listing of effective altruists working on helping animals.

The Strengths and Weaknesses of Different Cause Areas and Charity Ideas and **"EA" doesn't have a talent gap. Different causes have different gaps** by CE cover some of the basic differences between various cause areas. For animals, the key takeaway is that if you are running an evidence-based organization, getting funding from the EA community will be quite easy, but gathering evidence about cost-effectiveness will be much harder. Finding staff with the right skills may be more challenging as well.

Altruism, Numbers, and Factory Farms: This article lays out the general case for

why someone who is interested in effective altruism may focus on animal welfare. This is worth reading if you are undecided about whether or not to work on helping animals.

How to think about quantifying animal welfare

How Much Direct Suffering Is Caused by Various Animal Foods? By Brian Tomasik and follow-up post **How Much Suffering is in the Standard American Diet?** by Peter Hurford introduce some ways to think quantitatively about factory farming, introducing the framework of days-spend-in-factory per pound of product, and comparing the importance of different sorts of suffering (e.g. living conditions vs. pain of slaughter).

Is it better to be a wild rat or a factory farmed cow? A systematic method for comparing animal welfare is a post by CE working toward a cross-species system of “welfare points”.

Evidence gaps in the animal space (and their implications)

Global poverty could be more cost-effective than animal advocacy (even for non-speciesists) by Peter Hurford argues that the lack of evidence for cost-effectiveness in the animal space is severe enough such that GiveWell top charities may actually be helping more individual beings per marginal dollar than the most cost-effective animal charities. This means that even if you value humans and other species *equally*, if you have money to give you might actually still do the most good according to your values by donating to GiveWell top charities at the current time. This is a big part of why we are so enthusiastic about starting new animal charities – there’s a huge opportunity for improvement. Note that ACE had to cut **cost-effectiveness analyses** from their evaluation process altogether, perhaps in part because there just isn’t enough data to do an evaluation.

An opposing viewpoint is presented in **How good is The Humane League compared to the Against Malaria Foundation**. We think these comparisons are fraught because they do not **apply equal rigor** with respect to a wide variety of factors (to name a few such factors: less rigorous impact attribution, comparing marginal impact to total impact and not doing room for more funding analyses, lack of evidence discounting, etc). There are also many examples of people trying to calculate how much money you would need to save as many animals as one person going vegan (for example [here](#) and [here](#)) – we think these analyses are flawed for similar reasons and that it is very unclear to what extent additional donations to animal charities benefit animals

or avert animals suffering.

Please note: This should not be interpreted to mean that we think existing organizations are not doing good. We celebrate and admire their accomplishments, which have laid important foundations for our own thinking, and we think that they are doing important work bringing more talent and resources into the animal space. This is about the value of the *marginal donated dollar* in the animal space, and how it *compares* to the very high bar set by the most cost-effective charities identified by GiveWell within global health, which happens to be the most evidence-based field in the nonprofit sector and has a lot of infrastructure and resources devoted to gathering data. On the other hand, we think that additional *talent* directed toward these areas and organizations could be extremely beneficial. We encourage you to come to your own conclusion regarding this matter after reading up on [cost-effectiveness analysis](#), which will give you the tools and context to evaluate this for yourself.

Content by CE

Be sure to check out [Charity Entrepreneurship blog content](#). Information pertaining to animals can generally be found under the **ANIMAL RESEARCH** column, under the headings **Animal Research 2020**, **Process**, **Broad Research**, **Animal Research 2019**, and **Ask Research**.

Animals by the numbers

These are some of the most popular posts on the EA forum pertaining to counting animals. They are worth reading if you are working directly with any of the species named, or are trying to scout out [useful databases](#).

Accuracy issues in FAO animal numbers

[Estimates of global captive vertebrate numbers](#) by Rethink Priorities

[Rodents farmed for pet snake food](#) by Rethink Priorities

[Farmed Fish Welfare Report](#) by Animal Charity Evaluators. Notably, the report includes support for dissolved oxygenation, which is now being implemented by FWI.

[Fish used as live bait by recreational fishermen](#)

[35-150 billion fish are raised in captivity to be released into the wild every year](#) by Rethink Priorities

Analgesics for farm animals

Corporate campaigns and advocacy

These are the most popular posts on the EA forum specifically pertaining to corporate campaigns, and other forms of advocacy focusing on institutions. Worth reading if you are working on helping animals via institutional changes.

.

[Will companies meet their animal welfare commitments?](#) By Rethink Priorities.

[Corporate campaigns affect 9 to 120 years of chicken life per dollar spent](#) By Rethink Priorities

[35 Independent Pieces of Evidence for Why New Corporate Campaigns Might \(or Might Not\) Work](#) by CE

Careers and movement building

These are the most highly upvoted posts looking at jobs, skills, and networking within the EA animal community. Likely worth reading when thinking about making a hire, or if you are working on a meta-organization within the animal space.

[Effective animal advocacy movement building: a neglected opportunity?](#) by Jamie Harris, who later co-founded Animal Advocacy Careers with CE 2019 alumni Lauren Mee.

[Problem analysis of the talent bottleneck in EAA's](#) by Animal Advocacy Careers

[Effective Animal Advocacy Nonprofit Roles Spot-Check](#) by Animal Advocacy Careers

[There are *a bajillion* jobs working on plant-based foods right now](#)

Invertebrates and wild animals

Many of the most highly upvoted posts on the EA forum regarding animals involve moral interests regarding invertebrates and wild animals (most of whom are invertebrates). A few of our recommended charities involve wild animals and invertebrates. These are worth reading if you are working in a relevant area.

[Next Steps in Invertebrate Welfare, Part 1: Fundamental Research](#) By Rethink Priorities

[Next Steps in Invertebrate Welfare, Part 2: Possible Interventions](#) By Rethink Priorities

[Insect herbivores, life history and wild animal welfare](#) By Rethink Priorities

[The Importance of Wild-Animal Suffering](#)

[How Much Do Wild Animals Suffer? A Foundational Result on the Question is Wrong.](#)

Invertebrate sentience

Writing on the moral relevance of invertebrates. This is worth reading if you are undecided about whether you want to work with invertebrates, or if you need to prioritize between species in a way that involves invertebrates.

[Opinion: Estimating Invertebrate Sentience](#) by Rethink Priorities

[Invertebrate Welfare Cause Profile](#) by Rethink Priorities

[Invertebrate Sentience Table](#) by Rethink Priorities

[Invertebrate Sentience: A Useful Empirical Resource](#) by Rethink Priorities

[Invertebrate Sentience: Summary of findings, Part 1](#) by Rethink Priorities

[Invertebrate Sentience: Summary of findings, Part 2](#) by Rethink Priorities

[Why I'm focusing on invertebrate sentience – EA Forum](#)

General thoughts about sentience and philosophy

These may be worth reading if you need to prioritize across species, or are undecided regarding cause areas, or want to understand what other EAs are talking about philosophically.

[2017 Report on Consciousness and Moral Patienthood](#) by OpenPhil

[Comparisons of Capacity for Welfare and Moral Status Across Species](#) by Rethink Priorities

[Small animals have enormous brains for their size](#) – includes links to other interesting articles engaging in the brain size/neuron count/moral relevance discussion.

[Why I prioritize moral circle expansion over artificial intelligence alignment](#) – This isn't very practicality-oriented, but might be worth reading if you're frequently confused about terminology when EAs talk about things like "moral circle expansion" or "S-risk", or if you want to understand more about the perspective of long-termist leaning effective altruists who focus on animal rights.

3. The EA Community and Human Interventions

Author:

Ishaan Guptasarma,
Curriculum Developer & Content Writer

While effective altruism as a whole is interested in doing the most good in the abstract, there are many different ideas about how to achieve this in practice. Because of this, individuals within effective altruism often form subcommunities around “cause areas”, such as [animal advocacy](#), global health, [AI risk and other areas](#).

If you plan to create an organization within one of these cause areas and wish to recruit the support of the effective altruism community, it helps to have a good sense of what the community is talking about. We’ve compiled a series of articles that can help orient you to the current ways of thinking within each cause area.

This is not an exhaustive list of cause areas within EA, and we have prioritized writing about our focus areas for 2020. This is also not a list of all topics that we are considering prioritizing. For instance, even though we are recommending family planning interventions, family planning is not covered in this list, because outside of our own organization there has not been much discussion of family planning in the EA community (other than very philosophical questions of population ethics).

The articles that we most recommend you read are marked in **bold**. The articles not marked in bold were chosen less carefully, often on the basis of being popular on the EA forum, for being representative of how some segment of the community thinks, or for being one of the few examples of a post on a given topic.

Global health

This is the largest cause area within effective altruism. Most people who enter this program will already be familiar with EA involvement with global health in broad strokes, and most of our curriculum primarily references examples from this cause area. Most of EA’s ideas about how “cost-effectiveness” works are borrowed from

the larger global health community, so the differences between how EA thinks about this area and how the rest of the world thinks about this area are fairly small.

Our general recommendation for leveling up your knowledge of how EA relates to global health interventions is to make sure you know about the **EA organizations working in this space**. We also recommend you read almost **everything by GiveWell**, and keep any GiveWell **top charity** or **grant recipient** on your radar. They are an extremely effective and transparent organization, and you can learn a lot by looking at their methods. Of particular interest to charity entrepreneurs might be GiveWell's list of **charities we'd like to see**.

You can also check out our recommendations and research reports in our **blog** or use the "**poverty research**" filter.

Some other articles of general interest might include:

Lessons from an EA-aligned charity start-up (video) – Brendan Eappen on starting Fortify Health, and the differences between EA and the wider global health community.

Rob Mather: Against Malaria Foundation – What we do, How we do it, and the Challenges – Founder of a GiveWell top charity talks logistics and cost-effectiveness.

Worm wars: The fight tearing apart the global health community, explained: Figuring out how cost-effective deworming really is, is a significant issue in the broader world, and represents a key point of uncertainty for GiveWell as well.

Uncertainty and sensitivity analyses of GiveWell's cost-effectiveness analyses – This is the most upvoted EA forum post directly pertaining to strategy/epistemology strictly *within* randomista global health – that is, things that can be easily measured using randomized controlled trials such as vaccinations and mosquito nets.

Happiness and mental health

Mental health and happiness is one of CE's focus areas for 2020. The following is a collection of writing from the CE 2019 alumni who launched **Happier Lives Institute**. You can find more relevant content on the website itself.

- **Cause profile: mental health**
- **Is effective altruism overlooking human happiness and mental health? I argue it is.**
- Announcing the launch of the Happier Lives Institute
- Update from the Happier Lives Institute
- A Happiness Manifesto: Why and How Effective Altruism Should Rethink its Approach to Maximising Human Welfare
- Ineffective entrepreneurship: post-mortem of Hippo, the happiness app that never quite was
- Are You Sure You Want To Donate To The Against Malaria Foundation?

Why Charity Entrepreneurship is Researching Mental Health and Happiness

You can find more relevant CE content under the **Mental Health Research** heading on our [blog](#), or use the “mental health” tag

Mental health shallow review: Another independent review.

Policy and other institutional changes

Some of our recommendations this year are policy-focused interventions. Policy interventions differ from other types of interventions in a variety of important ways and tend to cluster together in terms of the practical issues involved in implementing and evaluating them.

These articles will not include policy interventions regarding climate change, as climate change is a separate section.

EAF’s ballot initiative doubled Zurich’s development aid: Describes how the Effective Altruism Foundation was able to increase Zurich’s development cooperation budget from \$3 million to around \$8 million per year, including a provision that the city will allocate it “based on the available scientific research on effectiveness and cost-effectiveness”. This is the first time Swiss legislation on development spending has explicitly mentioned cost-effectiveness.

How GiveWell’s research is evolving: Describes GiveWell’s intentions of engagement with some of the most cost-effective public health policy changes.

Growth and the case against randomista development is one of the most highly upvoted posts on the EA forum. It explains why lobbying for economic growth focused policies may be more cost-effective than randomista development.

Good Policies, which was launched via CE's 2019 incubation program, has published three articles: **Introducing Good Policies: A new charity promoting behaviour change interventions**, describing their launch and premise; **Tackling the Largest Cause of Death Worldwide: Good Policies Update on Tobacco Taxes**, describing what they've been up to lately; and **Why We Think Tobacco Tax Advocacy Could be More Cost-Effective than AMF**, which outlines their cost-effectiveness analysis.

OpenPhilanthropy is interested in several institutional and policy areas, including improving **Scientific Research**, **Macroeconomic Stabilization Policy**, **Biosecurity and Pandemic Preparedness**, **Immigration Policy**, **Land Use Reform**, and **Criminal Justice Reform** (in addition to the traditional EA focus areas of **animals**, **AI risk**, and **global health**).

An overview of **Biosecurity as an EA Cause Area**. (Note: EA interest in biosecurity is not directly related to and predates the recent COVID pandemic)

80,000 hours has career advice for **Improving institutional decision-making**.

Managing risk in the EA policy space – general advice, which was highly upvoted.

Thoughts on electoral reform – an outline for electoral reform as a cause area. Those interested in election science should check out **The Center for Election Science** as well as this list of **articles written by executive director Aaron Hamlin**. Aaron is one of the mentors of the Charity Entrepreneurship **Incubation Program**.

Increasing Access to Pain Relief in Developing Countries – a proposal for analgesic distribution that focuses on policy and institution level changes.

Which nuclear wars should worry us most? – one of several analyses by Rethink Priorities about nuclear policy

Existential Risk and Economic Growth – some long-termist discussion on whether economic growth is good or bad

Reading list

Here are some general resources containing items relevant to EA and human interventions not mentioned above.

[Introduction to Effective Altruism](#)

 [Major Funders within Effective Altruism](#)

 [Fundraising Outside of EA](#)

 [Key Donors in Global Health](#)

4. The EA Community and Other Interventions

Author:

Ishaan Guptasarma,
Curriculum Developer & Content Writer

Artificial intelligence

We're not planning on incubating organizations focusing on this area in 2020, but if you engage extensively with the EA community, it is something to know about. This is a cause area that applies to humans as well as animals. Interventions in this area tend to conceptually cluster with interventions within biosecurity, nuclear safety, and policy, involving a mix of advocacy and direct work on technical issues.

A general reading list:

The Sequences – Philosophical foundations. Useful for understanding various issues and improving epistemology in various ways, even if you don't work in AI. Also useful for understanding a lot of EA jargon.

AI and Effective Altruism – cause profile, relation to the movement

There's No Fire Alarm for Artificial General Intelligence – addressing critiques regarding lack of evidence

AI Alignment: Why It's Hard, and Where to Start – some technical explanations.

2019 AI Alignment Literature Review and Charity Comparison – a review of who is working on the problem and what they are doing.

Climate change

While most people who are interested in climate change within EA are primarily interested in helping humans, it impacts most creatures on earth in some way (though the impact on individual animals is unclear).

The EA community has relatively less engagement with climate change. Many EAs are interested in climate change, although some have argued that the space is already crowded, and that averting climate change is less cost-effective than traditional global health interventions (if you're near-termist) and isn't an existential risk (if you're long-termist).

[updated] Global development interventions are generally more effective than climate change interventions

Notes – Could climate change make Earth uninhabitable for humans?

Does climate change deserve more attention within EA?

Crowdfunding for Effective Climate Policy – a fundraising pitch for effective climate

policy, with interesting discussions in the comments.

Why we have over-rated Cool Earth – EA Forum

Extinguishing or preventing coal seam fires is a potential cause area

Founders Pledge Climate & Lifestyle Report

Climate Change Is, In General, Not An Existential Risk – EA Forum

Climate & ex risk

8 things I believe about climate change

Family planning

Family planning has not received much attention within effective altruism. We're pitching family planning as a cause area because we think that it not only measurably improves the lives of many individual women, but also brings systems-level change through advancing gender equality and various other positive ripple effects (e.g. [for the planet](#)).

For an overview and to learn more about our approach, we recommend you read **Why Is Charity Entrepreneurship Researching Family Planning**. We discuss benefits to the individual receiving access to family planning services with respect to mental health, physical health, education and income, and general women's empowerment. We also touch on additional effects in terms of animal welfare, climate change, and the effect on the total population.

Reproductive autonomy

How important do you think it is to empower people to have greater agency with respect to their bodies and futures, relative to other interventions you could be working on?

[Family Planning – Strategy overview](#) by the Bill and Melinda Gates Foundation

[Social and Economic Benefits of Reliable Contraception](#) – The Atlantic

[The Economic Effects of Contraceptive Access: A Review of the Evidence](#)

In addition to having significant positive effects on human health, happiness, and reproductive autonomy, we think that there are certain moral frameworks within which it can be a very high-impact intervention. The degree of impact differs depending on what one thinks about questions regarding the effects of family planning on various outcomes as discussed in the articles above, as well as on opinions regarding population ethics, impacts on nonhuman animals, and other effects.

Population ethics

Do you think creating new people is good, bad, or neutral? Your views on this subject might influence whether you want to work on family planning.

[What is the Demographic Transition Model?](#)

[How Important Is Population Ethics?](#)

[The Repugnant Conclusion](#)

Animal welfare

More people generally means more meat-eating, which means more animals in farms. This means that family planning might be incredibly cost-effective for animals. How should this weigh into our moral calculus, if at all?

[The Meat Eater Problem: Developing an EA Response – EA Forum](#)

[Quantifying the Impact of Economic Growth on Meat Consumption](#)

[How Much Direct Suffering Is Caused by Various Animal Foods?](#) by Brian Tomasik

[How Much Suffering is in the Standard American Diet?](#) By Peter Hurford, and see also our links to [databases](#) that quantify diet in other countries

Other

Here are some articles on the EA forum which are related to family planning as an intervention. (Note that there weren't too many posts on this topic to choose from in this category.)

[Women's Empowerment: Founders Pledge report and recommendations](#)

[The Effective Altruism Equality and Justice Project](#)

[Optimal population density: trading off the quality and quantity of welfare](#)

[Help in choosing good charities in specific domains](#)

5. The Strengths and Weaknesses of Different Cause Areas and Charity Ideas

Author:

Joey Savoie,
Co-founder & Director of Strategy

Unlike in previous years, we decided in 2020 to consider multiple different cause areas, which leaves more room for cause comparison. We think that generally, both entrepreneurs and donors have specific cause areas in mind when they attend or support our program. However, some have asked us for a sense of how the different cause areas, and more importantly, charities within them, compare.

We think each area has its strengths and weaknesses, and it's hard to reliably compare at this level because many assumptions (both ethical and epistemic) need to be made. This article offers a starting point for such comparisons. First we show roughly how our 2020 cause areas compare, and then we look back at and compare our recommended interventions from research done in 2019.

In 2019, our research focused on interventions for animals and within global health. As of 2020, we are considering the following four cause areas:

- Mental health
- Family planning
- Animals
- Health policy

2020 research: Comparing cause areas

The table below shows our **weighted factor model** framing for each 2020 cause area. Each area is color-coded from strongest to weakest.

Area	Direct cost-effectiveness	Relevant evidence	<u>Limiting factor*</u>	Execution difficulty	Non-captured externalities**
------	---------------------------	-------------------	-------------------------	----------------------	------------------------------

Mental health	Moderate	Some	Funding	Easier	EA movement
Family planning	Low	Moderate	Logistical	Complex	Child, animals
Animals	High	Low	Talent	Easier	Bar setting
Health policy	High	High	Funding	Complex	Precedent

* If the limiting factor cell is red, this means that the limiting factor will be met very quickly. Green means that the factor will be hard to meet.

** If the non-captured externality cell is green, this means that the externalities are large and positive. If the cell is red, this means that externality is small.

Another way to frame comparisons is in terms of more specific key strengths and weaknesses:

Area	Strengths	Weaknesses
Mental health	<p>Directness of the subjective well-being metric and possible underrating of the area by other metrics</p> <p>Possible promising cost-effectiveness for both low- and high-income countries</p> <p>Strong to moderate evidence base and background research but limited prioritization work</p> <p>Could encourage EA movement to consider more cause areas long term</p>	<p>Uncertain cost-effectiveness compared to top global health interventions</p> <p>More theoretical and philosophical work is required for assessment</p> <p>More limited funding base, particularly in the EA movement</p> <p>Evidence base has a wider range of metrics used, making it more difficult to compare</p>
Family planning	<p>Strong funding outside of EA</p> <p>Moderate evidence base</p> <p>Diverse range of positive effects (e.g. women's empowerment, unborn child benefits, family)</p>	<p>Maximizing multiple positive effects makes the charity harder to run</p> <p>Can be a controversial intervention</p>

	<p>benefits, income benefits, etc.)</p> <p>Under certain ethical views could be extremely impactful</p> <p>Area has more limited use of CEAs than others in global health, leaving promising subareas neglected</p>	<p>Size of impact depends in part on unsolved population ethics questions</p> <p>Evidence is spread out between a wide range of metrics thus speculative conversions and comparisons need to be used</p>
Animals	<p>Naive cost-effectiveness estimates generally show extremely high cost-effectiveness</p> <p>High levels of historical neglect mean many promising charity ideas are not yet founded</p> <p>Strong support both within and outside of the EA community</p> <p>Very strong case that animals should be given moral weight</p>	<p>Very low evidence base compared to other areas</p> <p>Some talent shortages in the movement that impair key charities</p> <p>More limited externalities and flow-through effects than other cause areas</p> <p>High rate of non-effectiveness minded activists in the area</p>
Health policy	<p>Naive cost-effectiveness estimates show higher cost-effectiveness than standard global health interventions and maybe all other human-focused areas</p> <p>Evidence base fairly strong if confidence is established in causal relationship of lobbying</p>	<p>Extremely complex space resulting in a much higher than average chance of a charity having limited or no impact</p> <p>More limited externalities and flow-through effects compared to other cause areas</p> <p>Very high bar of charities already working in the space</p>

2019 research: Comparing interventions

This section summarizes comparisons between the interventions recommended in our 2019 cause areas, global health and animals. As in the previous section, the two tables for each cause area show the weighted factor model and then intervention-specific strengths and weaknesses.

Global health interventions

Area	Direct cost-effectiveness	Relevant evidence	Limiting factor	Execution difficulty	Non-captured externalities
Immunization reminders	Low	Moderate	Funding	Easy	Limited
Tobacco taxation	High	Mixed	Policy windows	Complex	Precedent
Iron and folic acid fortification	Moderate	Moderate	Logistical	Moderate	Moderate

Area	Strengths	Weaknesses
Immunization reminders	<p>Highly flexible and has great feedback loops. For example can move from SMS reminders to SMS + gossip reminders</p> <p>Many impactful opportunities are still on the table due to the newness of the intervention</p> <p>Has a spectrum of success (unlike lobbying-based organizations)</p> <p>Easier to get government buy-in on large scale projects</p>	<p>High cost-effectiveness can be hard to achieve (e.g. higher than GiveDirectly but not as high as other GiveWell top charities)</p> <p>Behavior change makes concerns about external validity stronger</p> <p>Almost all benefits are in saving lives of children under five, so limited externalities</p>
Tobacco taxation	<p>Possibility of being extremely cost-effective</p> <p>Provides learning opportunities as well as the ability to set a precedent if success is achieved</p> <p>Often regarded as one of the most effective global health</p>	<p>High-risk opportunity with a high chance of failure</p> <p>Has active agents, namely tobacco companies, that push against interventions in this space</p> <p>The evidence base for the best</p>

	<p>policy interventions that is not consistently applied in all countries</p> <p>The evidence base that tobacco prevents a massive DALY burden and tobacco taxes reduce tobacco use is very strong</p>	<p>strategies to use to get tobacco taxes raised is unclear</p> <p>There are major difficulties in assessing the impact of a single organization in the space when multiple organizations are working in the same location</p>
Iron and folic acid fortification	<p>Nutrition as a broad area is seen as a highly promising area by a wide range of external experts including GiveWell and the Copenhagen Consensus</p> <p>Iron affects a wide range of health effects, often leading it to be undervalued in standard calculations. One example of this would be iron's effects on depression rates</p> <p>There are large gaps in fortification in lower-income countries</p>	<p>Certain locations are highly effective to run this intervention in (such as northern states in India) but there are fewer gaps than for other interventions</p> <p>Medium risk opportunity with a medium chance of failure</p> <p>There are many other nutrition-focused organizations, although none focusing on iron and folic acid in India</p> <p>Sub-standard fortification may lead to a limited or non-impactful effect</p>

Animal interventions

Area	Direct cost-effectiveness	Relevant evidence	Limiting factor	Execution difficulty
Dissolved oxygen for fish	Moderate	High	Logistical	Moderate
Food fortification for egg-laying hens	Low	High	Logistical	Low
Ask research	High	Moderate	Talent	Moderate
Animal careers	Moderate	Low	Replicability	Low

Area	Strengths	Weaknesses
Dissolved oxygen (DO) for fish	<p>Extremely cost-effective when compared to other animal-focused interventions. The most cost-effective direct intervention we measured</p> <p>Has a precedent in recent similar cage-free and broiler asks</p> <p>Strong evidence base</p> <p>Has strong pathways to funding as many donors consider fish a promising focus area</p> <p>Has the possibility of shifting the animal movement's fish focus in a much more effective direction</p>	<p>Upfront research required to determine key variables (e.g. optimal range of DO)</p> <p>Heterogeneity between species of fish makes it harder to generalize DO or other interventions</p> <p>Finding talent on the research side will be challenging in the animal space</p> <p>Fish focused charities would likely eventually get started so the counterfactual impact comes from the sub focus areas</p>
Food fortification for egg-laying hens	<p>Strong evidence base relative to other interventions in the animal space</p> <p>Less initial research needed before this organization could be founded</p> <p>Feed cost is the largest single item cost in poultry production</p> <p>Room for scaling to other food-related interventions</p> <p>Nutrition is a well understood and cost-effective intervention</p>	<p>The timing might not be optimal due to recent cage-free campaigns</p> <p>Lower cost-effectiveness than other animal charities</p> <p>Some concerns regarding counterfactual replaceability of the industry taking into account feed</p> <p>There is a wide range of possible nutritional improvements with exact effects on pain of birds being less clear</p>
Ask research	<p>There is limited research in both the animal movement as a whole and even less directly focused on asks that can be made of governments or corporations</p> <p>This type of research seems tractable and compared to other research, has quick feedback loops</p>	<p>The impact depends strongly on the effectiveness of corporate and governmental campaigns</p> <p>Impact relies on NGOs and organizations updating based on research</p> <p>Founders will have to be very strong in both research and communication skills</p>

	<p>Given the low cost of a research organization focused on this and that it is possible to affect large corporate campaigns, it could be highly cost-effective</p> <p>Effective altruists have a strong competitive advantage to found this idea</p>	<p>Relatively few asks are chosen annually, so the feedback loops are slow and there is downtime between key choices</p>
Animal careers	<p>Likely the charity idea the largest number of funders and activists will be excited about</p> <p>Meta-charity that could lead to other charities being founded</p> <p>Has a low floor for failure (even if done moderately well could have major benefits)</p> <p>Has models that can be replicated as a starting point (CSO in testing ideas, 80,000 hours in organizational scope)</p> <p>Could be very cost-effective if you take surveys of employee demand at face value</p> <p>Can be done in a wide range of locations effectively</p>	<p>Relies on other charities in the animal movement being net positive and effective</p> <p>Requires a high level of communication skills</p> <p>Requires a broad understanding of a diverse movement</p> <p>Co-founders will have to be comfortable with interorganizational interaction</p> <p>Many will not understand the charity idea or how it helps the movement</p> <p>Impact is indirect and hard to measure</p> <p>Very limited historical research in the area so starting from scratch in many cases</p>

6. Six Examples of New Charities in Effective Areas

Author:

Ishaan Guptasarma,
Curriculum Developer & Content Writer

In 2019, the Charity Entrepreneurship team graduated 13 alumni, who went on to launch six new charities, five of which implement interventions selected by our research program. How was this outcome accomplished, and how did these charities come into existence?

In 2016, Joey Savoie and Kat Woods started **Charity Science Health**, an organization that sends SMS text message vaccine reminders to people in India. Shortly afterward, they were approached by Brenden Eappan, a premed student who wanted advice on how to make an impact. Joey advised him to found a charity. He placed an ad for a co-founder, and was joined by Nikita Patel. The two of them started **Fortify Health** – an organization focused on getting major Indian wheat mills to supplement their product with iron and folic acid. Both **Charity Science Health** and **Fortify Health** went on to receive **GiveWell Incubation grants**. At time of writing, Fortify Health in particular is estimated to have a 25% chance of becoming a GiveWell Top Charity, and has been awarded a \$1,005,716 grant in order to scale up its operations. This was evidence that it could be high-impact to establish a full-time organization focused on identifying promising interventions and mentoring skilled individuals to execute them. As a result, Joey Savoie and Karolina Sarek decided to found Charity Entrepreneurship, a research and training program that incubates multiple high-impact charities every year.

The most important part of starting a successful charity is choosing a promising intervention in the first place. **Some interventions are much more cost-effective than others**, and even a very well-quantified, beautifully implemented program will not be the most cost-effective if the intervention itself is not cost-effective. It is common for researchers to have an in-depth knowledge of the particular intervention they are interested in. This makes a lot of sense if you are evaluating a single, preexisting intervention or charity. For Charity Science Health and Fortify Health, however, the aim was to implement one of the most promising of *all possible*

interventions. Therefore, the ideas for Charity Science Health’s “text message vaccine reminders” and Fortify Health’s “iron and folic acid supplementation of wheat” were chosen by spending thousands of research hours doing shallow reviews of approximately 45 potential intervention ideas, with a focus on doing *just* enough research to **eliminate non-viable candidates** and identify promising ones.

At Charity Entrepreneurship, this same methodology was used to launch the charities in the first formal Incubation Program, which took place in 2019. Using the above **research processes**, we generated **promising candidates** for good interventions, designed a **curriculum**, opened **applications**, and selected **promising potential founders**, as well as set up infrastructure to provide them with stipends, seed funding, and legal support. At the end of our program, six new charities were launched, which we introduce in this article.

Tobacco taxation

Our review found that bringing tobacco taxation policy in low- and middle-income countries in line with World Health Organization recommendations would save tens of millions of lives. The Center for Global Development calls tobacco taxation **the single best health policy in the world**, and GiveWell is investigating opportunities to influence government policy in this area.

To further this and similar policy-shaping efforts, we have helped Caleb Parikh and Joel Burke launch **Good Policies**, which is beginning work in both Mongolia and Armenia and is currently fundraising. Read their **Jan 2020 update on the EA forum**.

We’ve also helped Michal Trzesimiech (of **LEAN** and **EA Hub**) launch the **Policy Entrepreneurship Network**, aiming to identify unique opportunities in the space of health policy in low- and middle-income countries, as well as scalable ways of evaluating online campaigns.

Vaccination

A **randomized controlled trial** (RCT) by Nobel laureates Duflo and Banerjee suggests that word-of-mouth campaigns cost-effectively increase vaccination rates in India. To harness the power of nudges to combat vaccine-preventable diseases, we have helped Varsha Venugopal and Fiona Conlon launch **Suvita**. They are currently preparing to launch a pilot program to achieve operational proof-of-concept and to

better understand contextual factors that will guide iterative improvements.

Happiness research

We have provided incubation support to help Michael Plant and Clare Donaldson launch the **Happier Lives Institute** (HLI). HLI conducts theoretical research into the nature and measurement of subjective well-being (SWB), as well as applied research asking how resources can be best used to increase global well-being. They have recently hired economist Joel McGuire as a full-time research analyst, and are seeking **research interns** who can contribute to their **research agenda**.

Dissolved oxygen for fish

Our report finds that insufficiently oxygenated water is one of the main causes of suffering among farmed fish, which are numerous and historically have been neglected by the animal movement. We helped Haven King-Nobles and Thomas Billington launch the **Fish Welfare Initiative**, which aims to reduce the needless suffering of billions of fish by researching and executing targeted, highly scalable welfare interventions. After a preliminary survey of six different fish farms and processing facilities in four different countries, the Fish Welfare Initiative is conducting scoping research in South and Southeast Asia to identify the most promising regions in which to begin its pilot program.

Talent gaps in the animal movement

Many in the farm animal advocacy movement report that it is challenging to find the right staff to expand their organization. Staffing decisions are particularly difficult when an organization aims to expand into a new country or is hiring relatively senior staff. We've helped Lauren Mee and Jamie Harris launch **Animal Advocacy Careers**, where they have begun **systematic investigations of what works** in hiring and training in the animal movement. Their initial focus will be on addressing skills gaps in the movement, and Lauren has recently moved to Berlin to understand more about the challenges facing European organizations.

A Charity Entrepreneurship **Demo Day** gave each newly launched organization an opportunity to connect with mentors and funders in late 2019. All charities delivered lightning talks at **Effective Altruism Global 2019** in London. Each

organization also had a stand at the Effective Altruism Global Career fair, many looking to connect with potential hires within the effective altruism community.

We are awed and impressed by the hard work of each of these founders and wish them the best as they scale effective, impactful organizations. There is plenty left to do, and our organization will continue to incubate new charities every year. Not all the charity ideas we came up with last year were founded. There are still several equally promising ideas across a wide range of cause areas, and our team will continue to identify **more promising interventions**.

7. I Have a Charity Idea. Any Advice?

Author:

Kat Woods,
CE Co-founder

We often get asked for advice about a charity idea somebody has had. Every charity and entrepreneur will need different advice, but in this post we will cover the most cross-applicable advice that virtually everybody could benefit from:

- Generate and compare at least ten different options
- Really research your ideas – and try to destroy them.

Compare *at least* ten options

If there were one piece of advice I could give to practically anybody making any decision whatsoever, it would be to **compare more than two options**. Most people drift along in life, have an idea, then compare that single idea to the status quo. There are untold numbers of ideas out there; the odds are vanishingly low that any one idea is the optimal one. Furthermore, your choice can only be as good as the best option you have on the table. If there are only two, that puts hard limits on your potential upside.

To come up with a truly good idea, come up with at least ten different options. The “at least” is important. Ideally, you should try to think of dozens, then whittle them down. The more ideas you generate, the more likely it is you’ll come across something truly outstanding.

Starting a charity is a really big commitment, and the idea is the largest driver of impact. You could execute flawlessly, but if you’re running a charity distributing homeopathic medicine, you won’t be helping anybody. You want to make sure you’re committing to something worth the effort. Changing your direction later on is difficult, both logistically and psychologically.

When you brainstorm options, **don’t limit yourself to charity ideas**. You should also include alternatives such as working for a direct charity, earning to give, starting a for-profit, and any other thing you think might plausibly have an impact.

If you're feeling a bit apprehensive about such a complex decision, [here's a resource on what to do with all of your options once you've generated them](#) so as to make the decision more manageable. This is the process we follow, and you can learn [more about it on our blog](#).

Really research your idea – then try to destroy it

This may seem obvious, but most people I have spoken to have not really done this. Most people stumble into an idea, then ask a few people what they think. Those people don't want to be critical and "crush your dreams", so the person dives straight in.

If you want to really know if it's a good idea, research it. This is what EAs specialize in; play to your strengths! You can learn a lot about your idea before proceeding.

In fact, you can learn enough beforehand to know that you shouldn't actually do it in the first place. You should approach this with the attitude of trying to destroy your idea before you end up investing in and wasting time on something that you could have ruled out far faster.

Even more importantly, you want to rule it out before you become emotionally invested, or it's far less likely you'll ever be able to see that it's not a good idea. This is one of the greatest dangers lurking in the waters of charity entrepreneurship. If you are running a for-profit and it isn't making money, you'll eventually find out, no matter how much it hurts to admit. With a charity, the results are very rarely obvious. You could run something that has no impact, indefinitely, and fool both yourself and your donors. You could use your life on something that doesn't actually fulfill your values and never realize. While you [can set up systems to check and make sure you are indeed helping](#), it's easy to [ignore them, move the goalposts](#), rationalize the results away, or set up systems that don't actually falsify your program's impact, all in subtle ways that are hard for you to notice.

Of course, don't just try to destroy your program idea. Next, try to rebuild or strengthen the weaknesses you uncover. [Play steelman solitaire on it. Steelman](#) counterarguments then steelman countercounterarguments, and so forth. Maybe your idea won't work because you don't have experience in the area. How could you get around that? Maybe you could find a co-founder who does? Maybe you could do a few months of intensive study in a very small area? That there is a problem with

an idea does not necessarily mean that the problem is unsolvable. In fact problem solving is one of the most important skills an entrepreneur can have.

There are many potential weaknesses that your plan could have. One of the most common is finding out if it's been done in the past. Odds are somebody else has already had this idea and tried it. See if you can find them. Google it and ask people in area. If it's been done, learn as much as you can about what happened. If it didn't work, why not? Is it something you could plausibly solve, or should you abandon the idea? Is it already being run well, such that it might be better to simply support the existing program with money or time? Alternatively, if it's being run well but it's only running in a small section of its potential scope, you could use its success as a reason to scale it up, or run a similar program in another location.

For the research process in general, don't reinvent the wheel. We will be publishing material in the future which will provide a more in-depth account of how to choose an effective intervention, so subscribe to our [RSS feed](#) to learn more about this process. In the meantime, you can use [our research process as a](#) launching pad for your own process.

Another alternative is to [start a charity based on one of our recommended ideas](#). We have put in thousands of expert person-hours into the above process. Choosing them will give you a huge head start, alleviate analysis paralysis, and greatly increase the odds that your charity will be incredibly successful at making the world a better place.

In summary, if you're considering starting a charity, start with the fundamentals before thinking about how to fundraise and register as a not-for-profit. Figure out if your idea is a good one in the first place. Make sure to consider at least ten options and then research those, looking for disconfirming evidence before making the leap. Following these steps will greatly aid you in the challenging and exciting endeavor of starting a new, high-impact charity.

Part 2: Making Good Decisions

What is required to make good decisions? This is a complex question that would likely require several books to answer. Let's start with a simpler but related question: what tool is required to make a masterwork of wood? Is it the trusty hammer? Or perhaps instead the saw to cut away excess? Of course, the real answer is that it requires a large workshop of tools to make the most advanced works of wood.

Decision-making shares some significant similarities with woodwork. Just as you need hundreds of tools to make a masterwork in wood, you will have to employ a number of angles and techniques to make an optimal decision. This section of the handbook will go over a number of these tools and some of their applications.

To become a top decision-maker your work will never be done. There will always be more techniques to learn and master, but the ones presented here should get you started with some of the basics and get a sense of some of the most important tools. Think of it as the first toolset you buy when you move into your first own home. The set will not allow you to work on the trickiest of problems and you will acquire more tools as you need them, but it sure beats only having a hammer.

I. Meta tools (the toolbox)

Once you start collecting tools, the first thing you will notice is that you need some unified heuristics and techniques in terms of which tools to use when, and how long to use each tool for. Some techniques cross apply to a wide range of decision-making tools such as when to make a decision or how much time to give yourself to make a decision. This section covers some of these cross-cutting more meta techniques. They are the toolbox that allows you to store and organize your tools for optimal use. In this section, we will cover:

- When to focus and when to reevaluate
- How to determine how much time to spend on a decision
- How to narrow down a large number of options to the few ones to deeply consider
- How to use feedback to sharpen your tools
- Useful datasets that can help inform decision-making

1. When to Focus and When to Reevaluate


Author:

Joey Savoie,
Co-founder & Director of Strategy

We live in a complex world with ever-changing variables. When it comes to doing the highest impact activity, it can seem impossible to be confident in one plan over another. Additionally, new information from further research or even just time passing can make plan A seem better than plan B – only to have it switch again the next day. This can cause some people to be lost in **analysis paralysis** or to switch frequently between ideas, making little progress on each one. On the other hand, some people can pick a direction and never update again, constantly leaving impact on the table.

When I first entered the EA movement I was more the former, spending months on projects only to be convinced that something else was higher impact to work on

before the project could be efficiently finished. Not falling prey to sunk cost fallacy, I would switch to the higher impact project and start afresh. Sadly, ten half-finished projects do not equal the value of a single finished project. Although I was always working on what at the time I thought was highest impact, this lack of staying power made the overall output over those years fairly low relative to what it could have been if I had persevered through the projects. Of course, it's also very easy and common to get stuck or committed to a project that is no longer the best choice, and some people end up spending their career this way. So the real question becomes how to find a balance between those two. I think Sam Altman puts it quite well:



Although I think the specific numbers can be argued (perhaps it should be 6 months and 3 years) and for EAs generally our focus would be on the highest impact thing instead of the most interesting direction, I think the fundamental pattern is something I have also noticed. Using a strategy like this allows you to finish and mentally commit to large important projects. It also greatly increases the odds of a project coming to fruition. I expect a strategy like this will lead to far higher utility than a more intuitive strategy of “switch whenever you think something is higher impact than what you are doing” or “switch when something seems equally high impact but more exciting”.

Systematic exploring and choosing

Let's dive into a bit more detail about what both of these phases might look like. In the choosing phase you would want to be careful and fair-minded, likely doing systematic research on a broad range of options to determine the top contenders. You would not want to psychologically commit to any idea but remain impartial during this phase.

You likely will find many competing options that look promising, although different in their relative merits. You then want to compare them carefully and get external feedback from respected peers. This is the best time to get broad feedback as you are not committed to the project and can directly show the comparison. You could ask a close friend or mentor if they think plan A or B is better and explain the merits of both. Your friend could then answer honestly knowing that you have not already

made the decision. People often won't disagree with your plan if you're already committed, for fear of demotivating you or hurting your feelings.

The biggest enemy will be getting lost in analysis or being too quick to commit to a plan. Setting a specific amount of time for this research can greatly help both these concerns. **Time capping** is often necessary if you do want to make an endline choice because the truth is there will be very few occasions in project selection when you feel absolutely certain. Ultimately what you can hope for is to make a well-informed guess at the best option. And it's important to learn not to eliminate all uncertainty, but rather to **develop the art of being confident you made the best choice you could have in the time available**. On the other side of the coin, if you're predisposed to jump into the first good idea that comes along, seeing that you're only a fraction of a way through the time you set aside for research will remind you that you probably haven't carefully considered enough of the alternatives.

Relentless focus

A common topic in EA is the idea of donation splitting. A less discussed topic is the idea of how to split time between projects. The concept of how much focus to have or how many projects to split time between has applications both personally and organizationally. Many people will put careful analysis into where they donate (even if it's fairly small amounts of money) but not apply the same rigor to volunteering or generally spending time on projects. This is particularly true for projects that are alluring and seem to have minimal time commitments.

Two common approaches I have seen I will term the 90/10 approach and the heavy focus approach. In this context the 90/10 approach involves spreading one's time over a wide range of projects and getting large amounts of benefit from minimal amounts of time. For example, you might be on the board of many charities or have a full-time job plus several side projects. The heavy focus approach involves putting over 90%+ of your working time/energy into a single project (likely your full-time job unless you are a student or earning to give). An example of this might be a staff member of an EA organization who does not do any major projects on the side and instead puts more time/energy into working more hours for the organization.

Over time I have become more positive toward the heavy focus approach and less positive toward the 90/10 approach for most EAs, despite continuing to value it in many other topics such as research. One piece of evidence that updated me in this direction was the endline results of individuals who have been focused compared to individuals who have been more spread out. I have found that individuals who

spread themselves across many projects tend to have projects slip through the cracks and end up spending time on projects that after a comparative analysis would not compete with the top thing they are working on. When I compare these people to people of seemingly similar abilities but who have a much tighter focus, I tend to find that people in the latter category end up with much larger success and not a proportionally lower number of successes.

There are a few possible reasons this could happen. I think there is a heavy cognitive cost of switching your mind between very different projects, leading to less deep and creative thinking than you would expect given the split of hours. Generally, part-time workers have a harder time making major contributions to an organization and the same principle would seem to hold true between working more or fewer hours as a full-time employee.

A second possibility is that smaller projects tend to be hit harder by planning fallacy than large ones. I may expect it to take 3 years to start a high-impact charity, yet often it has ended up taking 5 years. But many small high-impact projects presumed to only take 10 hours ended up taking orders of magnitude more time.

Another possibility is that many of the most worthwhile projects are really difficult and require full-time attention to have a good chance of success. To use a personal example, most of my time is spent working on Charity Science Health, a direct poverty charity aiming to be GiveWell recommended. This project is sufficiently challenging that our odds are under 50% of being successful, even with almost all my time and energy going into the project. I think founding other charities or for-profits is often similarly challenging, and most of the very worthwhile projects also come with very high levels of difficulty and time commitment.

Plan reevaluation points

An important concept I often talk about to people starting new and ambitious projects is the idea of plan reevaluation points. With many ambitious projects, some days your project will be going great and you will feel great about it. Other days it will take all of your energy just to stay positive and everything will feel like it's falling apart. Days like this will come and go for even the best projects. It's great to think carefully about your project's impact but if as an EA you are constantly reevaluating your project, it will be hard not to abandon it on a hard day/week/month.

The solution I propose is to set “plan reevaluation points” where you review your project and plan with fresh eyes, looking over all the data you can to see if this project is still worth doing. I think how often to do these varies a bit depending on how long you have been working on the project. It might start as once every 6 months but move to longer periods until it’s once every 3 years. These plan reevaluation points can allow you a deeper level of focus without sliding into questioning your whole project, while still ensuring you update appropriately based on long term trends and evidence. The real trick to be able to hold off between plan reevaluations is having strong confidence in your initial choosing phase such that you know you made a good call. This system can save a lot of pain and also increase the chance you stay on your project long term.

In summary:

- Spend 1 year choosing and 4 years relentlessly pursuing a project.
- In the choosing phase, be careful, broad, systematic, and uncommitted. Pick the best option you can in the time you have given yourself.
- In the focus phase, avoid other projects that might seem like easy wins but will sap your time and energy. Avoid major project reevaluation until set points.

Following these steps, I think you can end up with a much higher impact career where you can make major progress on important projects.

This article was written in 2018 and [originally posted on the EA Forum](#).

2. The Importance of Time Capping

Author:

Joey Savoie,
Co-founder & Director of Strategy

Time capping can be defined as fixing the number of hours for a certain task, research project or decision, and keeping one's research within those bounds. Most tasks can be completed at different levels of depth and research itself is never-ending – a single topic could be researched in an hour, or could equally be the subject of an entire PhD. The same applies to website design, outreach, polishing or to the many other tasks that an organization engages in. Given tasks that are not time capped, people will generally spend more time doing what they find fun or intriguing, instead of what is best to put hours into in the long run. By setting a time cap on a task, we predetermine how important that task is relative to other counterfactual tasks. This approach often gets more done, albeit at some cost of depth – often 90% of the value of tasks is captured by the first 10% of the effort. There are a few specific kinds of tasks that benefit especially from time capping. Research and decision-making are tasks that frequently result in endlessly flexible deadlines without a clear 'good enough' point.

Research

Research is one of the foremost areas that can benefit from time capping. At Charity Entrepreneurship, we allocate a specific time frame for every research project. This has allowed our team to cover a lot of ground in a predictable way, and to cover the ground we need for making a specific decision within a clear time period (e.g. conducting research for a year and then recommending a list of top charities that could be started).

Many research teams could think of situations where twice as much research, but with half the time put into each piece, would likely result in more good for the world. For a more concrete example: we had a very specific time budget for our animal reports. For each of the reports, we could have gone much deeper (we ended up spending 1-5 hours per report and published a 1-page summary for 15 animals). But it would have been at the cost of breadth: for example, we could have spent 2-10 hours per report, but only covered 7 animals. Or, in theory, covered 30 animals at

half the depth. The key questions to consider were what the purpose of these reports was and how big a role they would play in our endline goal (i.e., starting effective charities). Essentially, what level of depth would give us enough information so that we could start to compare and consider which animals would be a priority? We decided that 1–5 hours would give us enough information for the soft prioritization of which animals to focus on.

I think highly intellectual cultures (such as the EA movement) tend to undervalue this sort of time-capped research approach, and I often see comments on research, even very deep research, that more or less translate to “put more time into this research”. Of course, sometimes it is true that more time should be put into a specific branch of research, but I very rarely see it the other way around – comments that roughly translate to “you should have put less time into this research”. It's become a cliché in academic research to say “**more research is needed,**” but in some areas this is really not the case, and particularly so once counterfactuals are taken into account. Time capping, I think, is the one strategy that can be used to improve this situation. If someone wants to make the case that we should have covered 7 animals, but at double the depth, I am very interested to hear the considerations in favor of making this trade-off, but if they say that more hours put into the research would make it better without a thought toward counterfactuals then it's harder to engage with.

Decision-making

Decisions are infinitely complex and, much like research, it would be easy for a person to spend anywhere between 1 minute and multiple years considering complex decisions. Some decisions are sufficiently complex that the perfect answer can never be found, just better and worse guesses. Given the uncertainty and complexity of the world, but also the importance of making many decisions (often hundreds a day), time capping the more lengthy decisions seems like an optimal solution. This stops analysis paralysis and constant indecisiveness. It also gives important decisions a fixed deadline and timeline for when they need to be made.

Using an example within CE: at the end of the year, we ultimately have to choose the best charities for us to recommend. This decision will almost by definition always be incomplete and subject to revision. But like with the research above, the important question is what are the counterfactuals? If CE, for example, did 5 years of research and decision-making in the animal space, would it result in higher impact charities recommended at the end? Almost definitely. But we also have to consider the impact of a strong charity (or multiple) that could have been founded 4 years earlier, as

well as the cost of research. CE could cover and recommend charities for mental health, poverty, the far future, and meta science. If two charities came from each of these areas during those four years, that would be eight charities not started for the end benefit of a better animal charity recommendation.

Some decisions are so complex that it's easier to have impact in multiple fields (or double the impact in a single 'best guess' field). Our team ended up thinking that it would be quicker and higher impact to run a year of CE on both animal and poverty issues than it would be to research and decide which one is of higher impact.

Using this system will not always result in the correct decision, but it makes up in time what it loses in precision. Sometimes the best decision is the one that is good enough, given limited time and knowledge. Time capping a decision puts into perspective how important this choice is in comparison to all the other choices, research, and tasks. It also allows you to look back on your decisions and feel as though you made the best choice you could in the time allocated for each.

3. Iterative Depth

Author:

Joey Savoie,
Co-founder & Director of Strategy

When making a decision or evaluating options, often hundreds of possibilities stretch as far as the eye can see. However, we often have limited time to explore these options and in many cases, it's impossible to deeply compare them all. A meta-process that can be used to narrow down from a large list to a much smaller one is the idea of **iterative depth**.

For example, the challenge of narrowing down from all the charity ideas in a cause area (often several thousand) to a reasonable number for deeper research can be a tricky process. A solution is to conduct multiple rounds of research (sometimes using different techniques) at different levels of depth. As the list slowly narrows down, the level of depth goes up. For our research process in 2020, that meant doing a quick **30-minute prioritization** of hundreds of ideas, then a longer **2-hour prioritization** of dozens of ideas, and finally an **80-hour prioritization** of the top 5-10. Each level of depth examines fewer ideas than the previous round but invests considerably more time into each idea.

This process can be cross-applied in most decision-making situations. If you are considering career options there are likely hundreds of options. But as you do a bit of research into each you will arrive at a shorter list, and maybe you consider **volunteering or interning** at the top 5. When considering different staff in hiring, you will go through a similar process of getting more information on fewer and fewer candidates until you have a final shortlist of a small number of candidates, but with multiple interviews and test tasks for each one.

A final example of using this process of iterative depth can be illustrated by a story of a charity that went through our program. They were working on a new and exciting idea in an area with lots of room to grow and test out different options. In fact, there were so many things to do that there was no way their organization would be able to cover all the ground. The organization's founder made a long list of

many actions they could take (several hundred). After reviewing the list they were able to sort it into rough tiers of promise, with about 20 ideas in the top tier. They then fleshed out these 20 ideas with a paragraph or two of explanation, and ran that list by a group of five trusted external advisors. The external advisors helped narrow down the list to three top actions to test out in the first year, with the plan to reevaluate later to determine whether it would be worth running more tests or scaling up one of the actions that worked best. This process helped narrow down to the true top options from what at first seemed like an impossibly large field.

Another way to frame iterative depth is that it is basically using systematic time capping on each stage of research. This ensures you don't end up having spent a large amount of time on a less promising option. Maybe you only have 100 hours to make a decision. Instead of spending all 100 hours across all possible options, you could break down your time. For example, you might spend the first two hours generating options, with a goal of coming up with 30 ideas. Then you could spend 15 hours spread across each idea (allowing 30 minutes per idea), and another 15 hours on cross-cutting research that could apply to a wide range of ideas. After that, you might narrow down the list to the top five options. Spend most of the remaining time split across those five options, say, 10 hours per option. Finally, you could spend the last hours comparing the two strongest options directly. After 100 hours you will have a good sense of which option seems best given the time you have assigned to it.

Iterative depth will not give you 100% confidence in the value of the ideas you ruled out. Yet it does give considerable transparency to your process. People can know how much time you spent on a given idea or how many options were considered before selecting your final one. It also tends to encourage brainstorming many different options, an invaluable meta-skill.

4. Useful Databases for Decision-making

Author:

Ishaan Guptasarma,
Curriculum Developer & Content Writer

Making good decisions requires accurate data about what is going on in the world around you. We use many databases in our research process to find candidates for the most cost-effective interventions.

For many of you, one of the first decisions you will make as an organization is **where to start your program**. Your choice of location can make a big difference to your chances of successfully overcoming initial logistical hurdles, as well as to your estimated cost-effectiveness in later stages. Additionally, once you pick a location, it can be difficult to change. Therefore, this is a decision worth spending a lot of research time on.

Making a good choice of location requires consulting real-world data. We've listed here some databases that we've found useful. This is, of course, not an exhaustive list. Additionally, these are domain-general databases, whereas the most useful databases to you will often be those that are specific to your intervention and (once you choose one) your location.

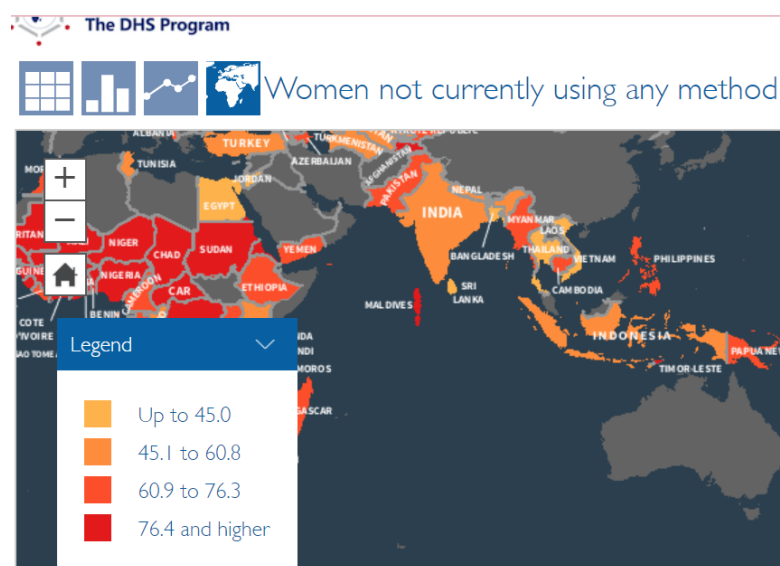
We recommend spending time searching for databases that suit your needs, as well as asking people who may have better knowledge of existing resources if they can point you toward useful information.

Databases for human interventions

General global health databases

The **Demographic Health Survey (DHS) Program** has a lot of useful data for global health (including mental health) interventions. You can use their **Spatial Database**

Repository to get detailed country, state, and local data on a wide variety of topics. If you'd like to work with a more visual map, **STATcompiler** has an especially user-friendly visual interface. You can, for example, select “all countries” and “women not currently using any method of contraception” to get a worldwide map depicting all surveys on that topic. Data is also sometimes available by subregion. There is a variety of information about non-health topics, such as mobile phone penetration and internet access.



The **Institute for Health Metrics and Evaluation** has not only a variety of useful health data, but also useful analyses and visualizations to help put the data in perspective and understand its effect on people. For example, the **Global Burden of Disease Study 2017** outlines the DALYs burden of various health problems by region. These databases played an important role in **Fortify Health's country/location selection**.




Image: Global DALY burden of various disease categories. IMHE features many **data visualization tools** such as this.

- The **World Bank** is a good source for quick country-level economic or health indicators. (Check out the World Development Indicators and the visualization [here](#).)
- The United Nations **Human Development Reports** host a variety of longitudinal indices, including the Human Development Index, the Global Multidimensional Poverty Index, and more.
- The **UN Data Hub** brings together data from all UN agencies.
- The **WHO Global Health Observatory** links to World Health Organization databases and indices.
- **Our World in Data** offers datasets spanning all areas from poverty to internet access.

Happiness and mental health

- The **World Happiness Report** provides data about happiness around the world.

Family planning

- The **Family Planning Database 2020** and the **UNFPA Family Planning Hub** specialize in data regarding fertility and contraception.
- The **World Population Prospect** specializes in population projections.

Logistics

- The **GSMA Global Internet and Phone Access** is one way to find out how easy it is to reach your beneficiaries.
- The **UK Foreign Office Travel Advice**, **US Travel Advisory**, **Crisis Group Crisis Watch**, and the **Armed Conflict Database** have information on safety and security for various locations.
- Consult the **World Bank Ease of Doing Business Rankings** to see how generally easy it is to operate in the country.

Other

- The **UNICEF Multiple Indicator Cluster Surveys** specialize in collecting data about women and children worldwide.
- **OECD Data** specializes in data regarding developed countries.
- The **IMF Data Hub** specializes in economic data and forecasting.

– Specific countries often have government websites with detailed information (e.g. [Nigerian National Bureau of Statistics](#)). There are also regional development banks (e.g. [Asian Development Bank](#)), and various open data repositories pertaining to specific locations or topics (e.g. [Open Data Portal for Africa](#) or [OpenAfrica](#))

Databases for Animals

The **Food and Agriculture Organization of the United Nations** (FAO) keeps data on food production to prevent malnutrition across the globe. This means that they also have various very detailed data that pertains to animal welfare, such as [what people eat around the globe](#), [where exactly different fish species are](#), and [FAOSTAT datasets](#) that allow you to ask very specific questions (e.g. how many beehives were farmed in Brazil in 2018). There are many useful datasets to be found, so be sure to click around the site. See also this post on [accuracy issues in FAO animal numbers](#).

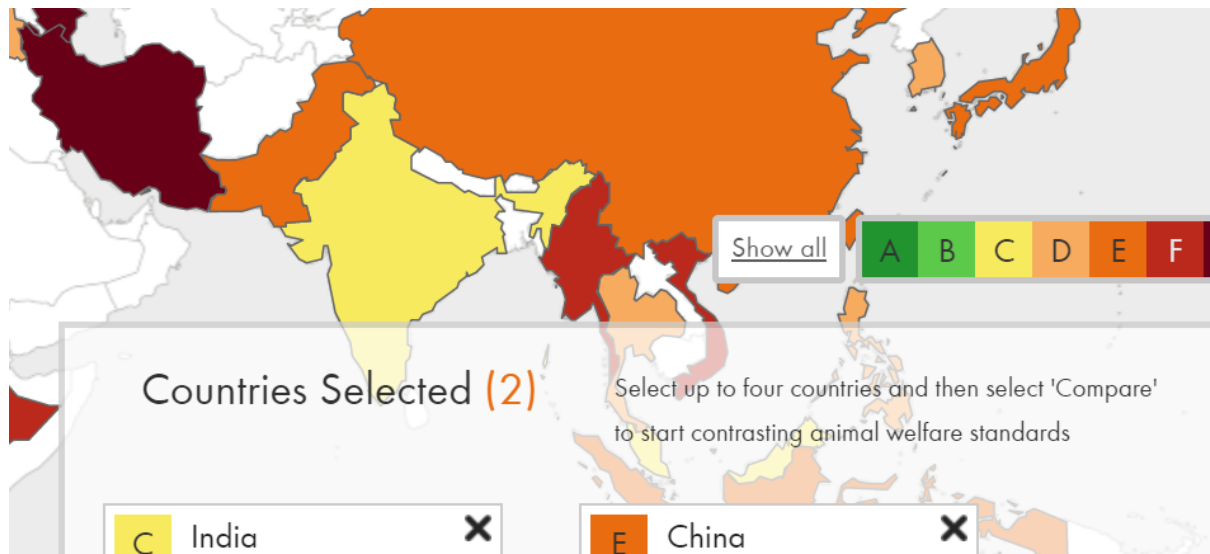


Image: *Inventory of food consumption surveys around the world*



Image: *Distribution of Atlantic Cod* from FAO's *FishStatJ* database.

The **Animal Protection Index** can help you compare the legal welfare standards for animals in various countries.



Animal specific databases: If you are working on an intervention for a particular animal, there are sometimes resources specific to that species. For example, Fish Welfare Initiative makes use of [FishEthoBase](#) and [Fishcount](#). Other specific databases include [EggTrack](#).

One good way to find databases is to go through research by other effective altruists focusing on animals and see which databases they cite for their research reports. See for example [Estimates of global captive vertebrate numbers](#) by Rethink Priorities.

Some other databases for animals include:

[European Commission on live animals](#) – track supply chains, data on disease

[Tridge](#) – imports, exports, production numbers

Our own [animal priority country spreadsheet](#) is a good place to dig around for resources as well.

[Estimates of global captive vertebrate numbers](#) by Rethink Priorities

Check out *Animals by the Numbers* under [The EA community and Animals](#) to see examples of others in the EA community running analysis.

Be sure to check out **Charity Entrepreneurship blog content** for articles that cite databases that you might use. Information pertaining to animal databases can generally be found under the **ANIMAL RESEARCH** column, under the headings **Animal Research 2020**, **Animal Research 2019**, and **Ask Research**.

Summary & resources

There's a lot of information out there, and you can often find data on questions that seem very specific and particular. Here's a condensed list of the databases we just reviewed:

Global health

Demographic Health Survey (DHS) Program (recommended)

Spatial Database Repository

STATcompiler

Institute for Health Metrics and Evaluation (recommended)

Global Burden of Disease Study 2017

Data Visualizations

World Happiness Report

United Nations **Human Development Reports**

World Bank

World Population Prospect

UN Data Hub

WHO Global Health Observatory

Our World in Data

Happiness and mental health

World Happiness Report

Family planning

Family Planning Database 2020

UNFPA Family Planning Hub

World Population Prospect

Logistics

GSMA Global Internet and Phone Access

UK Foreign Office Travel Advice

US Travel Advisory

Crisis Group Crisis Watch

Armed Conflict Database

World Bank Ease of Doing Business Rankings

Other

UNICEF Multiple Indicator Cluster Surveys

OECD Data

IMF Data Hub

Animals

Food and Agriculture Organization of the United Nations (FAO)

Inventory of food consumption surveys around the world

Where exactly different fish species are

Distribution of Atlantic Cod from FAO's **FishStatJ** database.

FAOSTAT datasets

Animal Protection Index (track legal protection for animals)

FishEthoBase

Fishcount

EggTrack (track corporate campaigns for egg laying chickens)

European Commission on live animals (track supply chains, data on disease)

Tridge (imports, exports, production numbers)

Estimates of global captive vertebrate numbers by **Rethink Priorities**

The EA community and Animals forum posts

Blog Content – CHARITY ENTREPRENEURSHIP

Example location selection

Fortify Health's country/location selection process (using databases)

5. Charities Love Feedback (As Long As It's Complimentary)

Author:

Joey Savoie,
Co-founder & Director of Strategy

Many charities claim to accept and even enjoy feedback (both positive and critical), but I find that the reality is less straightforward. Most charities have one of two kinds of set-ups for feedback:

1. Giving feedback is technically possible, but it is not encouraged, is inconvenient, or has negative ramifications.
2. Giving feedback is easy.

There is a huge difference between these two, but people often think both of these make an organization open to feedback.

Studies have shown that people are convenience maximizers, preferring to put in as little effort as possible. If something is a step harder, there is generally a huge drop-off rate. For example, if you ask someone to give feedback, but make it inconvenient or subtly discourage it, the amount of real feedback you get will be a small fraction of what would have been given if it was made easy.

If I wanted to make it seem like I was accepting feedback, but did not really want to receive or take into account large amounts of feedback (particularly negative feedback), I could do a number of things.

I could discourage feedback by:

- only accepting positive feedback;
- trivializing the feedback and making it seem unimportant;
- trying to censor negative public feedback.

I could make it inconvenient to give feedback by:

- asking others to only send me feedback privately;
- only accepting feedback from "chosen experts", who I personally pick;

- mainly putting my content on pages without the ability for others to comment;
- only telling the public about projects after their completion;
- responding to feedback very slowly so that the person giving feedback loses interest.

I could make it unpleasant to give feedback by:

- having a strong negative reaction (e.g. sadness, anger, defensiveness) to negative feedback;
- lashing out against those who give negative feedback (e.g. writing or saying negative things about what they are doing because of the feedback they have given);
- holding grudges against individuals or organizations that give negative feedback.

If I see an organization doing these harmful things with regard to feedback, I get the sense that they do not really want to improve based on others' suggestions. **Taking negative feedback in particular can be hard, but it's an important skill to learn and it allows ideas to be improved much faster.** We will never be able to solve the really important problems if we cannot admit that we are not doing everything perfectly.

II. Multi-factor decision-making (the Swiss Army knife)

If you were only allowed one tool, the Swiss Army knife would likely be your choice. By definition, it has multiple tools attached and can be useful in a large number of situations. Although a helpful tool, it would pale in comparison to using a more specialized tool or a full tool kit. The Swiss Army knife of decision-making is called multi-factor decision-making, and it involves combining multiple factors or variables into a single conclusion. Unlike the Swiss Army knife, this tool only gets stronger as you acquire other tools as its combining process can get more informed and sophisticated. This handbook primarily looks at multi-factor decision-making using spreadsheets, but although this is a convenient way to do it it's far from the only way. In this section, we will cover:

- Using a Spreadsheet to Make Decisions
- Why Location Matters When Picking an Intervention
- Geographic Selection Templates
- Assessments Made Easy With Google Spreadsheets
- How CE uses weighted factor models

1. Using a Spreadsheet to Make Decisions

Author:

Peter Hurford

We all make decisions every day. Some of these decisions are pretty inconsequential, such as what to have for an afternoon snack. Some of these decisions are quite consequential, such as where to live or what to dedicate the next year of your life to. Finding a way to make these decisions better is important.

The folks at Charity Science Health and I have been using the same method to make many of our major decisions for the past few years – everything from where to live to even deciding to create Charity Science Health. The method isn't particularly novel, but we definitely think it is quite underused.

Here it is, as a ten step process:

1. Come up with a well-defined goal.
2. Brainstorm many plausible solutions to achieve that goal.
3. Create criteria through which you will evaluate those solutions.
4. Create custom weights for the criteria.
5. Quickly use intuition to prioritize the solutions on the criteria so far (e.g., high, medium, and low)
6. Come up with research questions that would help you determine how well each solution fits the criteria
7. Use the research questions to do shallow research into the top ideas (you can review more ideas depending on how long the research takes per idea, how important the decision is, and/or how confident you are in your intuitions)
8. Use research to rerate and rerank the solutions
9. Pick the top ideas worth testing and do deeper research or MVP testing, as is applicable
10. Repeat steps 8 and 9 until sufficiently confident in a decision.

Which charity should I start?

For example, let's walk through the process that led to the founding of **Charity Science Health**. The definitive example for this process was **the Charity Entrepreneurship project**, where our team decided which charity would be the best possible charity to create.

Come up with a well-defined goal: I want to start an effective global poverty charity, where effective is taken to mean a low cost per life saved comparable to current GiveWell top charities.

Brainstorm many plausible solutions to achieve that goal: For this, we decided to start by looking at the intervention level. Since there are thousands of potential interventions, we placed a lot of emphasis on plausibly highly effective, and chose to look at **GiveWell's priority programs** plus a few that we thought were worthy additions.

Create criteria through which you will evaluate those solutions / create custom weights for the criteria: For this decision, we spent a full month of our six month project thinking through the criteria. We weighted criteria based on both importance and the expected variance that would occur between our options. We decided to strongly value cost-effectiveness, **flexibility**, and scalability. We

moderately valued strength of evidence, **metric focus**, and indirect effects. We weakly valued logistical possibility and other factors.

Come up with research questions that would help you determine how well each solution fits the criteria: We **came up with the following list of questions and research process**.

Use the research questions to do shallow research into the top ideas, use research to rerate and rerank the solutions: Since this choice was important and we were pretty uninformed about the different interventions, we did shallow research into all of the choices. We then produced the following spreadsheet:

	Program	Overall	Flexibility	Cost effectiveness	Metric focus	Logistical possibility	Indirect effects	Evidence	Scalability	Other factors
Malnutrition	Iron supplementation	8	Mid	Mid	High	Easy	High	High	High	Mid
Cash	Conditional cash transfers	8	High	Low-High	Low-High	Mid	High	High	High	High
Malnutrition	Super salt charity/multimicronutrient	8	Mid	High	Mid	Easy-mid	High	Mid	High	High
Behavior change	Mass media to promote behavior change	8	High	Low-High	Low-High	Easy-mid	High	Mid-low	High	High
Immunization	Immunizations	7	Mid-low	High	Mid-low	Hard	Mid	High	High	Mid
NCDs	NCD best buy	7	High	High	Low	Mid-Hard	High	Mid-low	High	Mid
Malnutrition	Treatment for severe acute malnutrition	7	Mid	Mid	Low	Easy	High	Mid-low	High	Mid
Malnutrition	Vitamin A supplementation	7	Mid	High	Mid-low	Easy	Low	High-mid	Low	Low
Data	Charities that collect or generate information and data relevant to GW	6	High	Unknown	Mid	Mid	Low	Low	Low	High
Malnutrition	Salt iodization	6	Mid	High	Lowest	Easy	High	High	Mid	Mid
Behavior change	Breastfeeding promotion	5.5	Mid	High	Low	Mid-Hard	Low	Mid	Mid-High	High
Malnutrition	Folic acid	5.5	Mid	Mid-low	High	Easy	Mid-low	High	Mid	Mid
Behavior change	Handwashing promotion	5	Mid	Mid-high	Mid-low	Hard	High	Mid-low	High	High
Targeted treatment	Pneumonia treatment	5	Mid-High	Mid	Low	Mid-Hard	High	Low	High	High
Diarrhea	Oral rehydration therapy	5	Mid-low	Mid-low	Low	Mid-Hard	Mid	High	High	Mid
Mass treatment	Deworming	4.25	Mid	High	Mid	Easy	High	Low	Mid-low	Mid
HIV	Provision of anti-retroviral therapy to treat HIV/AIDS	4	Low	Low	Low	Hard	High	Mid	High	Mid
Diarrhea	Therapeutic zinc supplementation	4	Mid-low	Mid-low	Low	Mid-Hard	High	Mid-low	High	Mid
HIV	Prevention of mother-to-child transmission of HIV	3	Mid-low	High	Low	Mid	Mid	Low	Low	Mid
Targeted treatment	Trachoma control (mass drug blind)	3	Mid-low	High	High	Hard	Mid	High	Mid-low	Mid
Targeted treatment	Tuberculosis case finding and first-line treatment	3	Mid-low	Mid	Mid-low	Hardest	Mid	High	Low	Mid
Cash	Unconditional cash transfers (external)	3	Mid	Low	High	Easy	Mid	High	Mid	Mid
Malaria	Malaria treatment	3	Low	Mid-low	Low	Hard	Low	Mid	Mid	Low
Mass treatment	Mass administration of ivermectin and albendazole to control lymphatic filariasis	2	Mid	Lowest	High	Mid	High	Low	Mid-high	Mid
Immunization	Immunization to prevent maternal and neonatal tetanus	2	Mid-low	High	Low	Mid-Hard	Low	High	Low	Low
Mass treatment	Mass administration of ivermectin to control onchocerciasis (river blindness)	2	Low	High	Mid	Mid-Hard	Low	Mid-low	Mid	Mid
Mass treatment	Polio eradication	1	Low	Unknown	Mid	Hard	Low	Mid	Very low	Low
Malaria	Bednet distribution	1	Low	High	Low	Easy	Mid	High	Mid	Low

Afterwards, **it was pretty easy to drop 22 out of the 30 possible choices** and go with a top eight (the eight that ranked 7 or higher on our scale).

Pick the top ideas worth testing and do deeper research or MVP testing, as is applicable / Repeat steps 8 and 9 until sufficiently confident in a decision: We then researched the top eight more deeply, with a keen idea to turn them into concrete charity ideas rather than amorphous interventions. When reranking, we came up with a top five, and wrote up more detailed reports – **SMS immunization reminders**, **tobacco taxation**, **iron and folic acid fortification**, **conditional cash transfers**, and a **poverty research organization**. A key aspect to this narrowing was also talking to relevant experts, which we wish we did earlier on in the process as it could quickly eliminate unpromising options.

Pick the top ideas worth testing and do deeper research or MVP testing, as is applicable: As we researched further, it became more clear that SMS immunization

reminders performed best on the criteria being highly cost-effective, with a high strength of evidence and easy testability. However, the other four finalists are also excellent opportunities and we strongly invite other teams to invest in creating charities in those four areas.

Which condo should I buy?

Come up with a well-defined goal: I want to buy a condo that is (a) a good place to live and (b) a reasonable investment.

Brainstorm many plausible solutions to achieve that goal: For this, I searched around on Zillow and found several candidate properties.

Create criteria through which you will evaluate those solutions: For this decision, I looked at the purchasing cost of the condo, the HOA fee, whether or not the condo had parking, the property tax, how much I could expect to rent the condo out, whether or not the condo had a balcony, whether or not the condo had a dishwasher, how bright the space was, how open the space was, how large the kitchen was, and Zillow's projection of future home value.

Create custom weights for the criteria: For this decision, I wanted to turn things roughly into a personal dollar value, where I could calculate the benefits minus the costs. The costs were the purchasing cost of the condo turned into a monthly mortgage payment, plus the annual HOA fee, plus the property tax. The benefits were the expected annual rent plus half of Zillow's expectation for how much the property would increase in value over the next year, to be a touch conservative. I also added some more arbitrary bonuses: +\$500 bonus if there was a dishwasher, +\$500 bonus if there was a balcony, and up to +\$1000 depending on how much I liked the size of the kitchen. I also added +\$3600 if there was a parking space, since the space could be rented out to others as I did not have a car. Solutions would be graded on benefits minus costs model.

Quickly use intuition to prioritize the solutions on the criteria so far: Ranking the properties was very straightforward. I could skip to plugging in numbers directly from the property data and the photos.

	Property	Mortgage	Annual fees	Annual increase	Annual rent	Bonuses	Total
A		\$7452	\$5244	\$2864	\$17400	+\$2000	+\$9568
B		\$8760	\$4680	\$1216	\$19200	+\$1000	+\$7976
C		\$9420	\$4488	\$1981	\$19200	+\$1200	+\$8473
D		\$8100	\$8400	\$2500	\$19200	+\$4100	+\$9300
E		\$6900	\$4600	\$1510	\$15000	+\$3600	+\$8610

Come up with research questions that would help you determine how well each solution fits the criteria: For this, the research was just to go visit the property and confirm the assessments.

Use the research questions to do shallow research into the top ideas, use research to rerate and rerank the solutions: Pretty easy, not much changed as I went to actually investigate.

Pick the top ideas worth testing and do deeper research or MVP testing, as is applicable: For this, I just ended up purchasing the highest ranking condo, which was a mostly straightforward process. Property A wins!

This is a good example of how easy it is to readapt the process and how you can weight criteria in nonlinear ways.

How should we fundraise?

Come up with a well-defined goal: I want to find the fundraising method with the best return on investment.

Brainstorm many plausible solutions to achieve that goal: For this, our Charity Science Outreach team conducted a literature review of fundraising methods and asked experts, creating a list of the 25 different fundraising ideas.

Create criteria through which you will evaluate those solutions/Create custom weights for the criteria: The criteria we used here was pretty similar to the criteria we later used for picking a charity – we valued ease of testing, the estimated return on investment, the strength of the evidence, and the scalability potential roughly equally.

Come up with research questions that would help you determine how well each solution fits the criteria: We created [this rubric with questions](#):

- What research says on it (e.g. expected fundraising ratios, success rates, necessary prerequisites)

- What are some relevant comparisons to similar fundraising approaches? How well do they work?
- What types/sizes of organizations is this type of fundraising best for?
- How common is this type of fundraising, in nonprofits generally and in similar nonprofits (global health)?
- How would we run a minimum cost experiment in this area?
- What is the expected time, cost, and outcome for the experiment?
- What is the expected value?
- What is the expected time cost to get the best time per \$ ratio (e.g., would we have to have 100 staff or a huge budget to make this effective)?
- What further research should be done if we were going to run this approach?

Use the research questions to do shallow research into the top ideas, use research to rerate and rerank the solutions: After reviewing, **we were able to narrow the 25 down to eight finalists**: legacy fundraising, online ads, door-to-door, niche marketing, events, networking, peer-to-peer fundraising, and grant writing.

Pick the top ideas worth testing and do deeper research or MVP testing, as is applicable: We did MVPs of all eight of the top ideas and eventually decided that three of the ideas were worth pursuing full-time: online ads, peer-to-peer fundraising, and legacy fundraising.

Who should we hire?

Come up with a well-defined goal: I want to hire the employee who will contribute the most to our organization.

Brainstorm many plausible solutions to achieve that goal: For this, we had the applicants who applied to our job ad.

Create criteria through which you will evaluate those solutions / Create custom weights for the criteria: We thought broadly about what good qualities a hire would have, and decided to heavily weight values fit and prior experience with the job, and then roughly equally value autonomy, communication skills, creative problem solving, the ability to break down tasks, and the ability to learn new skills.

Quickly use intuition to prioritize the solutions on the criteria so far: We started by ranking hires based on their resumes and written applications. (Note that to protect the anonymity of our applicants, the following information is fictional.)

	Person	Autonomy	Communication	Creativity	Break down	Learn new skills	Values fit	Prior experience
A	High	Medium	Low	Low	High	Medium	Low	
B	Medium	Medium	Medium	Medium	Medium	Medium	Low	
C	High	Medium	Medium	Low	High	Low	Medium	
D	Medium	Medium	Medium	High	Medium	Low	High	
E	Low	Medium	High	Medium	Medium	Low	Medium	

Come up with research questions that would help you determine how well each solution fits the criteria: The initial written application was already tailored toward this, but we designed a Skype interview to further rank our applicants.

Use the research questions to do shallow research into the top ideas, use research to rerate and rerank the solutions: After our Skype interviews, we reranked all the applicants.

	Person	Autonomy	Communication	Creativity	Break down	Learn new skills	Values fit	Prior experience
A	High	High	Low	Low	High	High	Low	
B	Medium	Medium	Medium	Medium	Low	Low	Low	
C	High	Medium	Low	High	High	Medium	Medium	
D	Medium	Low	Medium	High	Medium	Low	High	
E	Low	Medium	High	Medium	Medium	Low	Medium	

Pick the top ideas worth testing and do deeper research or MVP testing, as is applicable: While “MVP testing” may not be polite to extend to people, we do a form of MVP testing by only offering our applicants one month trials before converting to a permanent hire.

Which television show should we watch?

Come up with a well-defined goal: Our friend group wants to watch a new TV show together that we'd enjoy the most.

Brainstorm many plausible solutions to achieve that goal: We each submitted one TV show, which created our solution pool.

Create criteria through which you will evaluate those solutions/custom weights for the criteria: For this decision, the criteria was the enjoyment value of each participant, weighted equally.

Come up with research questions that would help you determine how well each solution fits the criteria: For this, we watched the first episode of each television show and then all ranked each one.

Pick the top ideas worth testing and do deeper research or MVP testing, as is applicable: We then watched the winning television show, which was **Black Mirror**. Fun!

Which statistics course should I take?

Come up with a well-defined goal: I want to learn as much statistics as fast as possible, without having the time to invest in taking every course.

Brainstorm many plausible solutions to achieve that goal: For this, we searched around on the internet and found ten online classes and three books.

Create criteria through which you will evaluate those solutions/custom weights for the criteria: For this decision, we heavily weighted breadth and time cost, weighted depth and monetary cost, and weakly weighted how interesting the course was and whether the course provided a tangible credential that could go on a résumé.

Quickly use intuition to prioritize the solutions on the criteria so far: By looking at the syllabi, table of contents, and reading around online, we came up with some initial rankings:

Name	Cost	Estimated hours	Depth score	Breadth score	How interesting	Credential level
Master Statistics with R	\$465	150	10	9	3	5
Probability and Statistics, Statistical Learning, Statistical Reasoning	\$0	150	8	10	4	2
Critically Evaluate Social Science Research and Analyze Results Using R	\$320	144	6	6	5	4
http://online.stanford.edu/Statistics_Medicine_CME_Summer_15	\$0	90	5	2	7	0
Berkley stats 20 and 21	\$0	60	6	5	6	0
Statistical Reasoning for Public Health	\$0	40	5	2	4	2
Khan stats	\$0	20	1	4	6	0
Introduction to R for Data Science	\$0	8	3	1	5	1
Against All Odds	\$0	5	1	2	10	0
Hans Rosling doc on stats	\$0	1	1	1	11	0
Berkeley Math	\$0	60	6	5	6	0
OpenIntro Statistics	\$0	25	5	5	2	0
Discovering Statistics Using R by Andy Field	\$25	50	7	3	3	0
Naked-Statistics by Charles Wheelan	\$17	20	2	4	8	0

Come up with research questions that would help you determine how well each solution fits the criteria: For this, the best we could do would be to do a little bit from each of our top class choices, while avoiding purchasing the expensive ones unless free ones did not meet our criteria.

Pick the top ideas worth testing and do deeper research or MVP testing, as is applicable: Only the first three felt deep enough. Only one of them was free, but we were luckily able to find a way to audit the two expensive classes. After a review of all three, we ended up going with “Master Statistics with R”.

2. Why Location Matters When Picking an Intervention

Author:

Ishaan Guptasarma,
Curriculum Developer & Content Writer

If choosing the right intervention is one of the most important parts of starting a highly effective charity, then choosing the correct location to run that intervention is often the second or third big important choice made by new founders. Your choice of location can make a big difference to your chances of successfully overcoming initial logistical hurdles, as well as to your estimated cost-effectiveness in later stages.

Once you pick a location, it can be difficult to change. Aside from general expenses, many location-specific networks and skills may not transfer from place to place. Therefore, this is a decision worth spending a lot of research time on.

Finding the greatest addressable need

When working on global health interventions focused on humans, it's usually best to work in areas with a large population that suffers from severe poverty, and especially locations that suffer from easily preventable health issues. How many potential beneficiaries stand to benefit from your intervention? A proportionately smaller number of people who are in need live in more affluent locations such as New York, London, or Tokyo, and it is generally more difficult to alleviate those types of suffering using a very small amount of money.

Cause	Organization	Cost	Impact
Developing-world health	<u>Malaria Consortium (seasonal malaria chemoprevention program)</u>	Approximately \$2,300 per life saved ²	Improve health, save lives
Early childhood care and education (U.S.)	<u>Nurse-Family Partnership</u>	\$15,000 per child served ³	Increase academic performance and reduce criminal behavior
US Education	<u>KIPP</u>	\$9,000-\$20,000 per student per year (including state funds) ⁴	Improve academic performance

Image source: *Your Dollar Goes Further Overseas*. All three of these charities are highly regarded and are doing very good work. They represent very cost-effective organizations within their respective domains. Nevertheless, due in part to location considerations, a donation to Malaria Consortium ultimately accomplishes more good. Organizations working in areas with a high burden of very harmful but easily preventable issues are able to use money to do good more cost-effectively.

It's often also helpful to work in locations that present fewer logistical challenges. Locations such as Syria, South Sudan, Afghanistan, or the D.R.C. have many people who are suffering from issues such as malnutrition and disease, which may at face value seem easily preventable. However, it is often difficult to alleviate this suffering in a cost-effective manner because violence, political instability, and lack of infrastructure make it difficult to operate in these countries (though certainly possible – for instance, AMF has net distributions in both Sudan and the D.R.C.).

Finally, it's important to work in a population where there's a large number of potential beneficiaries for your intervention. Before pivoting to using conditional cash transfers for vaccinations, New Incentives used cash transfers to incentivize HIV-positive pregnant women to give birth in a clinic, where appropriate procedures can minimize the risk of disease transmission to the infant. While the intervention was cost-effective, the organization was unable to scale because there wasn't a large enough population of potential beneficiaries – and they were spread out across hundreds of clinics.

Working in locations such as New Delhi has the advantage of juxtaposing the infrastructure of a city (for instance medical supplies, helpful government agencies, qualified staff) and close access to extreme poverty. These are also locations that have a large population and high population density, ensuring a large pool of potential beneficiaries. In these regions, you will often find situations where you can accomplish a great deal of good for many people at low cost. Populations in such locations will often be served by other organizations, and it's worth keeping track of who else is working in this space.

Remote regions are often poorer than large cities. Even these areas are usually relatively easily accessible from nearby towns (within a few hours' drive). Northwest Nigeria, where New Incentives operates, has several larger towns that function well as operational hubs and Kano is a megacity with access to international flights.

This advice can change quite a bit depending on the particulars of what one is working on. When working on policy interventions, for example, things like population size and disease burden often still matter, but the degree of affluence might be less important and it might be cost-effective to work in middle- or perhaps even high-income countries. When working with animals, it's still worth thinking about the total number of animals involved and the severity of the welfare issues in various locations, but there are new logistical considerations to consider. Many logistic issues overlap (e.g. safety issues), but some are different: for example, two locations may have a similar number of animals, but convincing a single corporation to implement a change in practice across their production facilities is different from convincing many unregulated small scale farmers to do the same. It's good to make a list of which considerations are most likely to apply to your intervention.

How location influences logistical hurdles

Running a highly cost-effective intervention generally means you must make use of local resources rather than building everything from scratch. If you want to distribute vaccines, you don't open an entire hospital – you work with an existing hospital to subsidize or incentivize treatment. If you want to fortify food, you don't open a mill, you work with existing mills and enable them to fortify. Therefore, existing infrastructure makes a big difference regarding how you will interact with your beneficiaries. Some common factors that might make a difference include:

- General infrastructure: This depends on the intervention you are working on, but it's tough to work with hospitals, schools, and other institutions that are chronically understaffed. This problem can be fairly severe in some cases – for example, some cash transfer programs are focused on paying health workers more to incentivize them to show up to work.
- The “informal sector”: In low-income countries, it's common for the majority of the population to work in the “informal sector”. That means that they are starting businesses or being employed without any formal contracts with or recognition from any government. Informal supply chains have many implications for interventions based on lobbying and corporate campaigns.

- Access to financial services: Any intervention that requires exchanging money with beneficiaries becomes easier if they have access to some form of financial service.
- Phone access: Many interventions require you to be in contact with your beneficiaries. Also note that having a phone means that beneficiaries may have access to mobile financial instruments (for example M-Pesa).
- Access to electricity: Low or inconsistent access to electricity spells trouble for many interventions. For example, it might be difficult to keep medicine refrigerated, or to keep aerators for fish running (many fish farms are just holes dug into the ground, nowhere near electricity). This also means that beneficiaries will often have discharged phones and therefore be uncontactable.
- Proximity to talent: If you want to hire skilled people who live near your intervention site and who share a language in common with you, it's often easier if the intervention site is near a large city where a lot of qualified people might live.

In addition to logistical hurdles relating specifically to interacting with your beneficiaries, there are also logistical hurdles to consider regarding the day-to-day realities of running an organization.

- General lack of standardization: Some places use multiple languages, with dialects changing sharply from place to place. You often can't rely on buses and trains being consistent or predictable. In some places, you can't take for granted the idea of "addresses", which can make finding beneficiaries tough. People might not use the same calendar as you, or any calendar. Legal standards are often inconsistently applied as well.
- Corruption: In some places, it is difficult to do anything without being asked to pay bribes. This can happen even when dealing with law enforcement, government officials, and other authorities one might go to for recourse.
- Bureaucracy: Some places have more bureaucratic systems that are less streamlined and more difficult to navigate. You may find yourself struggling to do things such as open a bank account or open a phone line.
- Health and safety: You and your staff may hesitate to work near certain places due to health and safety concerns. Some factors which may be detrimental to health and safety might include low availability of medical care, local diseases, food safety, availability of drinking water, high traffic conditions, air pollution, targeted harassment due to race, gender, or sexual orientation, military conflicts, and crime.

Not all of these factors will matter to you, and there will likely be additional factors not listed here that will. When choosing a location, do some thinking about the **logistics of the intervention you picked** and try to get data on the factors which are likely to be most relevant to you.

How location influences cost-effectiveness estimates

(Recommended reading: **Using more rigorous CEAs: New Incentives, Fortify Health, and CSH**)

If you are able to overcome the initial logistical hurdles and set up your charity, the next challenge is to achieve a level of cost-effectiveness that is sufficient to justify scaling. Location makes a big difference when it comes to these factors.

When I attended Charity Entrepreneurship's 2019 incubation program, one of my "practice tasks" was creating a CEA for incentivizing vaccinations in Delhi.

I remembered that **GiveWell had done a CEA of New Incentives**, which gives cash to people who get vaccinated in the northwestern region of Nigeria.

New Incentives Sheet	Median	Andrew	Chelsea	Chris	Elie
calculation	results	subjective input	input to be updated	by	more data
Summary & Results					
New Incentives vs AMF	0.78	0.77	1.24	0.74	0.54
New Incentives vs GiveDirectly	5.15	4.30	14.12	10.89	4.90
Cost per equivalent outcome to an under-5 death averted	\$1,963	\$1,854	\$1,707	\$2,871	\$2,490
Percent benefit driven by preventing under-5 deaths		95%	98%	98%	97%
Percent benefit driven by consumption benefits		5%	2%	2%	3%
PCV (vaccine against pneumonia)					
Vaccination rate without the program: PCV 1		37%	33%	33%	37%
Vaccination rate without the program: PCV 2		30%	29%	29%	30%
Vaccination rate without the program: PCV 3		23%	24%	24%	23%
Vaccination rate without the program: exactly 0 doses of PCV		63%	67%	67%	63%
Vaccination rate without the program: exactly 1 dose of PCV		8%	4%	4%	8%
Vaccination rate without the program: exactly 2 doses of PCV		7%	4%	4%	7%
Vaccination rate without the program: exactly 3 doses of PCV		23%	24%	24%	23%
Vaccination rate with the program: PCV 1		52%	53%	43%	47%
Vaccination rate with the program: PCV 2		45%	49%	39%	40%
Vaccination rate with the program: PCV 3		38%	44%	34%	33%
Vaccination rate with the program: exactly 0 doses of PCV		48%	47%	57%	53%
Vaccination rate with the program: exactly 1 doses of PCV		8%	4%	4%	8%

I also remembered that GiveWell had done a CEA of Charity Science Health, which uses SMS text messages to remind people when it is time to get vaccinated in India. (I say “India” because the CEA used country-wide data. This isn’t ideal, of course, since India is large and non-homogenous. In reality, Charity Science Health operated out of New Delhi and covered Uttar Pradesh, Madhya Pradesh, Rajasthan, Delhi NCR region, Maharashtra, West Bengal, and Bihar – but these were not analyzed individually by the CEA.)

Charity Science Health	Median	Andrew	Chelsea	Chris	Elie
Summary & Results					
Charity Science Health vs AMF	1.0	1.0	1.4	1.3	0.9
Charity Science Health vs GiveDirectly	7.1	5.4	16.0	19.0	8.1
Cost per equivalent outcome to an under-5 death averted	\$1,500	\$1,482	\$1,503	\$1,643	\$1,503
Percent benefit driven by preventing under-5 deaths		100%	100%	100%	100%
PCV (vaccine against pneumonia)					
Vaccination rate without the program: PCV 1		90%	90%	90%	90%
Vaccination rate without the program: PCV 2		89%	89%	89%	89%
Vaccination rate without the program: PCV 3		87%	87%	87%	87%
Vaccination rate without the program: exactly 0 doses of PCV		10%	10%	10%	10%
Vaccination rate without the program: exactly 1 dose of PCV		2%	2%	2%	2%
Vaccination rate without the program: exactly 2 doses of PCV		2%	2%	2%	2%
Vaccination rate without the program: exactly 3 doses of PCV		87%	87%	87%	87%
Vaccination rate with the program: PCV 1		94%	94%	94%	94%
Vaccination rate with the program: PCV 2		93%	93%	93%	93%
Vaccination rate with the program: PCV 3		92%	92%	92%	92%
Vaccination rate with the program: exactly 0 doses of PCV		6%	6%	6%	6%
Vaccination rate with the program: exactly 1 doses of PCV		1%	1%	1%	1%
Vaccination rate with the program: exactly 2 doses of PCV		1%	1%	1%	1%

I decided that the best way to figure out what a cash transfer program in India might look like was to use GiveWell's New Incentives CEA, but import *country data* from GiveWell's Charity Science Health CEA. If you haven't already, you should read [Using more rigorous CEAs: New Incentives, Fortify Health, and CSH](#) before proceeding with this section. The remainder of this section will assume a general familiarity with these CEAs.

Here is a table summarizing the differences between Charity Science Health and Fortify Health. As you can see, within a cohort of 100,000 beneficiaries, Charity Science Health vaccinates far fewer people – but is also much cheaper (due to operating in India and not using cash transfers), and ultimately comes out as more cost-effective in this analysis.

	GiveWell Staff Median Rating	
	New Incentives	CS-Health
Summary & Results		
versus AMF (depends only slightly on moral weights since AMF also primarily saves under-5s)	0.7825240249	1.0
versus GiveDirectly (depends strongly on moral weights on consumption vs. lives saved)	5.15	7.1
Cost per equivalent outcome to an under-5 death averted (median rating)	1962.864078	\$1,500
Benefits		
Under-5 deaths averted by the program	73.04116716	4.33
Total consumption benefit for 10,000 person cohort $\ln(\text{income} + \text{cash transfer}) - \ln(\text{income})$	80.00510817	
Moral weight conversion from "consumption" to "under 5 death averted"	2.312647399	
Total benefit (equivalent under-5 death averted), per 10000 person cohort	75.35381455	4.33
Costs		
Cost of cash transfers per 10,000 person cohort	57391.9624	
Administrative costs	88478.18182	
Costs to the government of additional vaccines and clinic load per 10,000 person cohort	7351.792116	
TOTAL COST	152906.2269	\$6,500
This value is (median benefit)/(median cost), unlike A7 which is median(benefit/cost)		
Cost per equivalent outcome to an under-5 death averted (median benefit/median cost)	2029.176994	\$1,501

We want to take the Indian country data from CSH and apply it to the New Incentives CEA, to estimate what would happen if we ran a similar cash transfer program in India.

This has been done in the spreadsheet [New Incentives modified with CSH country data](#). I've marked all items that have been imported from CSH, as well as all my explanatory notes, in **bold, 14pt font** to make it easier to find. How was this spreadsheet done?

First, we need to recalculate the costs.

Costs	Costs re-estimated using Numbeo index			
Cost of cash transfers per 10,000 person cohort	57391.9624	35200.4036	Delhi cost of living	18.86
Administrative costs	88478.18182	54266.61818	Lagos cost of living	30.75
Costs to the government of additional vaccines and clinic load per 10,000 person cohort	7351.792116	4509.099164	Delhi:lagos ratio	0.6133333333
TOTAL COST	152906.2269	93782.48585		

For this, I used Numbeo's [cost of living plus rent index](#), and used the ratio between the index value in Delhi and the index value in Lagos to estimate how the costs might change in Delhi. Delhi is cheaper to live in than Lagos, so they're quite a bit lower.

Then, we need to recalculate the benefits. We'll take some of the [disease & vaccine parameter calculations](#) from the CSH CEA and port them to New Incentives.

All country-specific data is ported

Pertussis (disease)		
Yearly chance of infection, if susceptible	10%	CSH CEA estimates used
Weekly chance of infection, if susceptible	0.20%	
Case fatality rate, under 12mo	3.7%	
Case fatality rate, 1-4yo	1.0%	

Whereas the yearly chance of being infected with Pertussis in Northwestern Nigeria is 30%, in India it is only 10%. This is obviously an important consideration when thinking about the relative importance of providing vaccinations for this disease to Northwestern Nigerians vs. Indians (but bear in mind that much of this reflects the choice to use statistics for “India” as a whole – a more localized analysis would reveal that different Indian states have wildly different health outcomes).

Some items represent assumptions that are held constant.

Disease & Vaccine Parameter Calculations		
Average age at Visit 1, in weeks, without program		2
Average age at Visit 2, in weeks, without program		6
Average age at Visit 3, in weeks, without program		10
Average age at Visit 4, in weeks, without program		14
Average age at Visit 5, in weeks, without program		39
Average age at Visit 5, in weeks, without program		39

We don’t need to change everything. Some items, such as the average age at each visit, are simply assumptions that are held constant. Ideally, these would be updated with empirical data, but it seems that this data was not available for either charity and so an estimate from a different source was used.

Some medical statistics are assumed constant (reason unknown)

Tuberculosis (disease)		
Tuberculosis: Chance of infection if susceptible, yearly	1.15%	location independent estimates (the numbers for CSH are the same)
Tuberculosis: Chance of infection if susceptible, weekly	0.02%	
Tuberculosis: Case fatality rate	27%	
Tuberculosis: Uncertainty adjustment for case fatality rate	100%	
Tuberculosis: Adjusted case fatality rate	27%	
Proportion of live neonates who have been infected by and survived tub	0.00%	

Some items remain constant between the two spreadsheets because the two sections held the same value for every cell. Here, the same risk-of-infection and rates-of-fatality if infected were used for estimating tuberculosis statistics in both Nigeria and India. I’m not sure why GiveWell kept these constant. It might be that they were estimated from studies in such a way that we hope they generalize across location. It is also possible that there was no sufficiently trustworthy estimate of this type available for India. In reality, of course, these things are not always constant from place to place. For example, obviously all these numbers would be non-negligibly different if we were talking about beneficiaries in Germany.

Some things were kept the same because I didn't have time to change them correctly

Routine MCV1 (Measles vaccine)					
Vaccination rate without the program: MCV1	32%	30%	26%	87%	Wasn't changed to CSH data (It ought to be, but doing it correctly would be too complex)
Vaccination rate with the program: MCV1	42%	43%	37%	92%	
Vaccine efficacy at preventing disease when administered around 9mo	73%	73%	73%	77%	
Uncertainty adjustment for MCV1 vaccine efficacy	100%	100%	100%	100%	If you want to try, the CSH data is in the "anyone" column.
Adjusted MCV1 vaccine efficacy	73%	73%	73%	77%	

If New Incentives (using a cash incentive) increases measles vaccination from 30% to 43% in Nigeria, and CSH (using a reminder) increases the vaccination rate from 87% to 92% in India, how much should we model a cash incentive increasing the vaccination rate in India? **If this were a real CEA, this would be very important, and not something to leave as is.** One should come up with clever ways to estimate this, or look for studies that try to estimate this. However, because that would require a lot of time and this is just an illustrative model for education, I've left the New Incentives values as they are.

Results: Summary table

	GiveWell Staff Median Rating			
	New Incentives	CS-Health	New Incentives CEA with Indian data	
Summary & Results				
versus AMF (depends only slightly on moral weights since AMF also primarily saves under-5s)	0.7825240249	1.0		Readme: These numbers are the result of taking New Incentive's Lagos CEA, swapping out the location data parameters (e.g. baseline vaccination rates, mortality rates) with Charity ScienceHealth's CEA in India, and then adjusting for cost of living differences between Delhi and Lagos.
versus GiveDirectly (depends strongly on moral weights on consumption vs. lives saved)	5.15	7.1		
Cost per equivalent outcome to an under-5 death averted (median rating)	1962.864078	\$1,500		
Benefits				
Under-5 deaths averted by the program	73.04116716	4.33	27.67819039	
Total consumption benefit for 10,000 person cohort Ln(income+cash transfer)- Ln(income)	80.00510817		80.82218635	
Moral weight conversion from "consumption" to "under 5 death averted"	2.312647399		2.338594572	
Total benefit (equivalent under-5 death averted), per 10000 person cohort	75.35381455	4.33	53.5219606	
Costs			Costs re-estimated using Numbeo index	
Cost of cash transfers per 10,000 person cohort	57391.9624		35200.4036	Delhi cost of living 18.86
Administrative costs	88478.18182		54266.61818	Lagos cost of living 30.75
Costs to the government of additional vaccines and clinic load per 10,000 person cohort	7351.792116		4509.099164	Delhi:lagos ratio 0.6133333333
TOTAL COST	152906.2269	\$6,500	93782.48585	
This value is (median benefit)/(median cost), unlike A7 which is median(benefit/cost)				
Cost per equivalent outcome to an under-5 death averted (median benefit/median cost)	2029.176994	\$1,501	1752.22441	

As you can see, the cost-effectiveness of New Incentives changed quite a bit when some of India's country data was put in. The benefits reduced (because India as a whole has a better baseline vaccination rate than Northwestern Nigeria) but the costs reduced as well (because Delhi is cheaper than Lagos), putting the final cost-effectiveness of our hypothetical charity somewhere midway between New Incentives and Charity Science Health. While it didn't happen in this analysis, sometimes, differences in cost-effectiveness between locations can be orders of magnitudes different from each other.

Of course, bear in mind that this number would likely be better if we used data from a more specifically targeted region in India, and it would likely be worse if we had more properly accounted for the degree to which incentives might be less effective at improving vaccination rates in a context where the baseline vaccination rate is higher, rather than leaving it undone because it was too complicated. It's also important to note that these CEAs are outdated, and quite a bit of new data has come out since they were done which would change these cost-effectiveness estimates. However, the overall point is that features of the geographic location in which you choose to do your intervention make a huge difference. If you pick a location with factors that are favorable to your cost-effectiveness, you can have a much bigger impact.

Summary

Choosing a location is one of the most important decisions you will make about your intervention. Logistical hurdles differ from place to place, and choosing a place with many hurdles can make the initial setup and scaling of your charity challenging.

As we've demonstrated using CEAs from New Incentives and Charity Science Health, location considerations can dramatically alter the ultimate cost-effectiveness of your charity when the time comes to scale. Making this choice correctly can be the difference between success and failure.

In the next article, we will be showing you how to use a spreadsheet to make decisions about where your location should be, combining both logistical and cost-effectiveness considerations.

Resources

Location and logistics

[How logistics will influence which intervention you pick](#)

Location and CEA

[Using more rigorous CEAs: New Incentives, Fortify Health, and CSH](#)

[New Incentives CEA Winter 2017](#)

[\[Early 2017\] Charity Science Health CEA](#)

[New incentives modified with CSH country data](#)

3. Geographic Selection Templates

Author:

Ishaan Guptasarma,
Curriculum Developer & Content Writer

In the previous article, we explained why location matters in terms of logistics and cost-effectiveness. This section will provide more specific templates to help you make your own location decisions, by explaining the process that other organizations associated with Charity Entrepreneurship have used to make decisions about their location.

If you're unfamiliar with the concept of using spreadsheets to make decisions, start with [Using a Spreadsheet to Make Good Decisions: Five Examples](#). If you're unfamiliar with navigating Google Sheets in general, consider reading [The Ultimate Guide to Google Sheets](#) by Zapier. You'll also find useful information in our chapter on [Assessments Made Easy with Google Spreadsheets](#).

Choosing a location – Fortify Health

Fortify Health's [Country/location selection](#) decision spreadsheet includes the disability-affected life years (DALYs) burden of anemia (remedied by iron), neural tube defects (remedied by folic acid), and malaria risk (because iron supplementation may increase malaria risk) in each potential location. It also estimates important logistical factors, such as centralization of wheat supply chains, cooperativeness and effectiveness of the government, and safety; as well as personal factors particular to the founders, such as language barriers and preexisting networks of helpful contacts. Finally, there are estimates regarding crowdedness in the field, with regard to supplementation programs that already exist.

metric	DALYs Rate (/100,000 population)	DALYs Rate	DALYs Rate	Prevalence	Deaths	Deaths	Deaths	Population (millions)	DALYs	Greatest untapped potential of centralized fortification (g/person/day produced in industrial mills not currently fortified)	Political barriers and conflict (these were rapid subjective ratings initially used for country selection - see adjacent columns with worldbank data for improved objectivity)	These data (presented as z-scores) corroborate subjective scoring used in initial selection, but were added to this sheet post hoc for comparison			population (log scaled)
view as	Iron-deficiency anemia	Neural tube defects	Malaria	Iron-deficiency anemia	Iron-deficiency anemia	Neural tube defects	Malaria	People	IDA + NTDs			Government Effectiveness	Political Stability and Absence of Violence/Terrorism		
condition															
Average of top 20 after first												1.06	-0.67	-0.48	0.39
Average of top 5 after first												1.20	-0.51	-0.22	0.31
Filter/sort in this row -->															
India	1235	26	248	38%	0.03	0.06	3.65	1324.2	1261	94.8	0	0.1	-0.955	0.99	
Tanzania	697	125	2279	24%	0.06	1.07	29.42	55.6	822	19.8	0	-0.55	-0.488	0.56	
Mauritius	655	25	0	27%	0.01	0.08	0.01	1.3	680	296.4	0	0.958	1.048	0.04	
Chad	1336	671	3574	37%	0.04	7.31	44.32	14.5	2007	0.0	2	-1.494	-1.208	0.37	
Mauritania	918	354	509	30%	0.03	3.61	6.44	4.3	1272	118.0	2	-0.787	-0.736	0.20	
Zambia	561	88	2440	18%	0.07	0.65	31.28	16.6	649	113.8	2	-0.658	0.183	0.39	
Ghana	922	237	5133	28%	0.70	2.27	69.67	28.2	1159	4.0	2	-0.202	-0.162	0.46	
Ethiopia	472	85	118	17%	0.07	0.55	1.31	102.4	552	47.3	2	-0.639	-1.569	0.64	

Fortify Health's **State Selection** decision spreadsheet goes into greater detail regarding consumption and supply chains of *atta*. It narrows down which states have a confluence of anemia and high consumption of *atta* from centralized supply chains (such that it would be easy to introduce and monitor fortification in just a few mills and have a large impact).

	State total population (millions)	Women with any anemia	Children with any anemia	Men with any anemia	Ratio of women with anemia to men with anemia	Current wheat flour fortification efforts known to FSSAI in government schemes or open market as of 3/2018	Percent of accounted for consumption centrally produced (CCM+RM)	Percent of accounted for consumption that is centrally produced excluding maids ATTA ONLY (CCM+RM)	Percent of all ATTA only that is centrally produced (CCM+RM)	Times per capita total wheat consumption centrally produced as atta (millions) <proxy for total reachable population>	Times per capita total wheat consumption centrally produced as atta per 10 mills (millions) <proxy for reachable population per 10 mills>	State total wheat consumption accounted for in 2011-2012 FFI supply chain analysis	% of surveyed consumption accounted for across market channels in FFI analysis
Maharashtra	112.4	48.0%	53.8%	17.6%	2.7	3 districts MDM flour, 2 districts PDS flour	43%	23%	29%	16.5	2.4	3,340,000 (64%)	
West Bengal	91.3	62.5%	54.2%	30.30%	2.1	PDS flour is thought to be fortified outside	91%	32%	79%	30.5	3.1	1,707,900 (104%)	
Rajasthan	68.6	46.8%	60.3%	17.2%	2.7	NA	31%	18%	21%	10.4	0.8	5,648,490 (83%)	
Jammu & Kashmir	12.5	49.4%	54.5%	20.7%	2.4	1 brand (NavBharat-Vitamin Plus) registered	89%	42%	79%	6.2	1.3	553,266 (120%)	
Jharkhand	33.0	65.2%	69.9%	29.9%	2.2	NA	77%	42%	65%	2.6	0.4	208,880 (19%)	
Uttarakhand -	10.1	45.2%	59.8%	15.5%	2.9	NA	43%	30%	35%	3.3	0.2	799,761 (107%)	

Choosing a location – Charity Science Health

Charity Science Health's **State Prioritization List** includes birth and vaccination rates, as well as data regarding mobile penetration and literacy. This is because the intervention primarily involves increasing vaccination coverage for children during their first years of life, and the mechanism to accomplish this is sending text messages reminding parents to take their children to update their vaccinations when it is time. The intervention thus requires that the targeted population have a mobile phone to receive the message, and literacy to be able to read the message.

Entire State D2D Ranking- Program Effect	State	NFHS 2015-16 All basic vaccination rate	RSOC Full Vaccination Rate 2013-2014 -	Average Vaccination Rate((NFHS+RSOC)/2)	Mobile Penetration rates (NFHS 2015-16)	Literacy rates NFHS 2015-16	Thousands of Birth Score	Thousand births per year
2.5	Uttar Pradesh	51.1%	68.6%	59.9%	91.8%	56%	1.00	5460
2	Gujarat	50.4%	56.2%	53.3%	90.4%	65%	1.00	1274
2.5	Madhya Pradesh	53.6%	53.5%	53.6%	84.0%	52%	1.00	1931
2	Assam	47.1%	55.3%	51.2%	87.5%	65%	1.00	701
3	Bihar	61.7%	60.4%	61.1%	90.3%	45%	1.00	2875
3.25	Jharkhand	61.9%	64.9%	63.4%	83.7%	55%	1.00	814
2.5	Rajasthan	54.8%	60.7%	57.8%	93.6%	50%	1.00	1777
3.5	Andhra Pradesh	65.3%	74.1%	69.7%	88.5%	56%	1.00	864
2.75	Maharashtra	56.3%	77.4%	66.9%	88.9%	74%	1.00	1865
2.5	Uttarakhand	57.7%	68.6%	63.2%	94.6%	72%	0.37	187
3.25	Jammu and Kashmir	75.1%	59.0%	67.1%	95.9%	63%	0.44	221
3.5	Telangana	68.1%	--	68.1%	89.0%	57%	1.00	618

Exercise: Are there any important considerations that might be missing from Charity Science Health's spreadsheet? Are there any items from Fortify Health's spreadsheet that might be applicable to Charity Science Health?

Choosing a location – Animals

Charity Entrepreneurship has some **preliminary country research** that should be generally helpful for animal-focused organizations. This includes data on farmed meat production by species in number of individuals, consumption by species in weight, legal protections for animals, trade agreements, public attitudes, English proficiency, and other useful metrics.

Country	Overall	Funding/Producti on	Current funding	Overall production (individual s)	Production: Fish (MM)	Production: Hens (individuals)	Production: Turkey (individuals)	Tractability /45
China	1	\$0.0000160	\$1,108,611	69142279100	62172687500	1638181818	0	low
India	1	\$0.0002269	\$3,102,830	13676078388	12376518000	310505000	0	medium
Indonesia		\$0.0000430	\$500,000	11639898063	9475874000	133176764	0	medium
Bangladesh	1	\$0.0000023	\$10,000	4430314441	4026168000	8015500	0	medium
Vietnam	1	\$0.0000023	\$10,000	4283873286	3896144000	47141000	0	medium
Egypt	2	\$0.0000031	\$10,000	3252429027	3069809396	34302600	1703000	low

Fish Welfare Initiative shared that they used the following methodology to make this research work for them:

- Use the spreadsheet to generate a longlist of promising countries
 - Pick countries that do well on the spreadsheet, as well as countries that experts or respected sources recommended.
 - Put in some other countries, especially those that are unusual in some way (to test the validity of internal metrics).
- Create a middle-list

- Spend 30 min reading each CE report (found in the **Country Rating** subs spreadsheet) and 15 minutes writing up pros and cons for each item on the long list.
 - Eliminate half of the least promising options to create a middle-list.
- Create a shortlist
 - Spend 2 hours identifying crucial considerations which might help narrow things down, and then researching them or asking experts. **Crucial considerations** are concepts that might radically shift your estimates about how likely something is to work. They might include things like: “Do interventions involving oxygenators work on a farm with inconsistent electricity?” or “How centralized is the production of fish in each region?”
 - Eliminate the least promising option.
- Create a top 5
 - For each remaining item, spend 2 hours reading and talking to experts until you can identify the top 5.
- Rank the top 5. Visit the top 5 countries if it helps narrow it down. Spend 12 hours writing up reports based on the findings.

Summary

Once you have identified the location-driven factors that are most important for your success, both in terms of logistics and cost-effectiveness, you are ready to begin researching which location is best suited for you to operate in. Look at these examples of how some of our other incubatees have chosen their location to create a research process of your own.

In brief, you need to optimize for the following:

- Cost-effectiveness. This usually means low baseline rates of e.g. vaccinations, and high poverty, which results in higher mortality.
- Logistics and scalability. This usually means high population numbers as a first step and ease of scalability as a second step (e.g. mobile phone penetration rates for a vaccination reminder charity, government openness to your program, safety, etc.)

Resources

Templates from other organizations

[Charity Science Health State Prioritization List](#)

[Fortify Health Country Selection](#)

[Fortify Health State Selection](#)

Data and other resources

[Useful Databases for Decision Making](#)

[Priority Country Spreadsheet](#) (for Animal interventions)

Resources for using Google Sheets effectively

[Using a Spreadsheet to Make Good Decisions: Five Examples.](#)

[The Ultimate Guide to Google Sheets](#) by Zapier

[Assessments Made Easy with Google Spreadsheets](#)

4. Assessments Made Easy With Google Spreadsheets

Author:

Ishaan Guptasarma,
Curriculum Developer & Content Writer

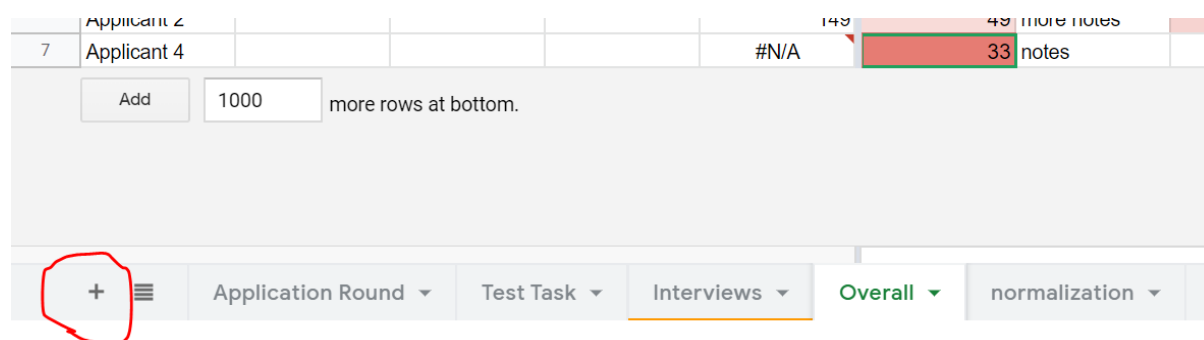
This section presumes some knowledge regarding the use of Google Sheets. Here is a quick knowledge check to see if you would benefit from this section:

Are you generally familiar with the method of ranking things on a spreadsheet via various criteria? If not, you may want to read [Using Spreadsheets to Make Decisions](#) first.

Can you use [sumproduct](#) to quickly add scores and multiply them by weight, e.g. `=sumproduct(F3:K3, F$2:K$2)`? Would you know how to automatically color-code cells according to what was written in them? If you answered no to either of these, you may want to read [Zapier's Guide to Spreadsheets](#) first.

Getting all your data in one place

Integrating data from multiple spreadsheets can be tricky. Luckily, there are a few tricks you can use. One helpful hint is to keep every phase of your interview process in a single Google Sheet. You can add multiple spreadsheets (also known as tabs) in a single document by clicking on the '+' on the bottom right of the screen.



You can pull data from other spreadsheets that are in the same document. For example, if you set a cell `=Interviews!D5` it sets the value to cell D5 of the sheet named "Interviews". You can also just navigate to a cell, press '=', and then

navigate to a different part of the sheet in order to pull values from the cells that you want. All of the ordinary functionalities of Google Sheets (e.g. Ctrl-D, functions, etc.) will apply across multiple spreadsheets.

For example, if you are doing a hiring round and have your candidates neatly sorted, it is sometimes pretty easy to manually copy-paste or directly refer to things in order to put all the data in one place. However, there may come a time when your data is too large to manually search through. If you don't want to go through a spreadsheet by hand to pull up all this data, one especially useful tool for pulling data is the vlookup, or “vertical lookup”, function. You can use this to search an entire spreadsheet for a string of text, and then pull data from a specific column of that row.

For example, suppose I have an applicant named “Applicant 1” and I want to pull up the notes I took on their resume, but I don't want to manually search for their name. I know that the notes are somewhere in the Application Round spreadsheet, on column 7.

I can use `=VLOOKUP(A3, 'Application Round'!$1:7, 4, FALSE)` where A3 is the cell containing the name of the applicant, 'Application Round'!\$1:7 specifies the spreadsheet and domain I want to search in (you can just click over the region you want and the text will auto-generate). Now it will search the entire domain of 'Application Round'!\$1:7 for a row containing the text of A3 (which happens to be “Applicant 1”) and return whatever is in column 4 of that row.

The screenshot shows a Google Sheet interface. The formula bar at the top displays `=VLOOKUP(A3, 'Application Round'!$1:7, 4, FALSE)`. Below the formula bar, a table is visible with the following structure:

	A	B	C	D	E	F	G	H	I	J
1	Name	Notes	Offered?	Accepted	Overall additive score	Initial application score	Notes	Test task	Notes	Interview
2					Weights	1		3		
3	Applicant 1		Yes	Pending response	142	57	<code>=VLOOKUP(A3, 'Application Round'!\$1:7, 4, FALSE)</code>			
4			Pending info on				great fit but			

I can use the Ctrl-D function to automatically do this for *all* my applicants! (Just be careful to use the \$ functionality to freeze appropriate reference points when using Ctrl-D – in this example, add the \$ to `Application Round'!$1:7`).


Interpreting the data

Remember that spreadsheet scores are only a tool to aid your thinking. At this stage, if you only look at numbers, it's very easy to trick yourself.


Take a look at these scores. They were generated via a straightforward use of `=sumproduct` on raw scores, weighting the test task highest and the initial application score lowest. If you just look at the overall score, it looks like applicant 1 is the worst candidate.

Name	Notes	Offered?	Accepted	Overall additive score	Initial application score	Notes	Test task	Notes	Interview	notes
				Weights	1		3		2	
Applicant 1		Yes	Pending response	142	57	Application notes	13	notes for task	23	interview notes
Applicant 5		Pending info on logistical issues		159	65	great fit but logistical issues	14	notes for task	26	interview notes 3
Applicant 3				#N/A	55	notes notes	7	rejected for task	#N/A	#N/A
Applicant 2				149	49	more notes	10	notes	35	interview notes 2
Applicant 4				#N/A	33	notes	#N/A	#N/A	#N/A	#N/A

But if you look at the component bits, you can see that Applicant 1 is only *slightly* behind the *best* candidate on the test task (1 point) and the interview (3 points)! That could easily be a matter of chance. Additionally, Applicant 2, who is in second place, did pretty poorly on everything except the interview.



What's going on? The problem here is that we're adding metrics that were calculated using very different processes. Even though we only *weighted* the "initial application score" a "1", we awarded points in the range of 33–57 for the application – a 24 point spread. Meanwhile, even though the test task is *weighted* 3, it was awarded between 7 and 14 points, which translates to 21–42 points after adding the multiplier – a 35 point spread. Finally, for our interview, we awarded it 23–35 points, weighted 46–70 points, which is a 24 point spread. Even though we *wanted* to weigh the test task highly and the application process very little, the application process ended up contributing a lot to the process because it was calculated using high numbers and contributed a lot of variability.



Because each of our metrics used a different measurement system, the initial application process (blue) ended up being weighted much more highly than we actually intended, which unfairly penalized Applicant 1. This extra weighting doesn't reflect importance, only the coincidence that the score tended to use bigger numbers.

At this point, we might not want to go back and change the entire grading rubric, so we will have to get creative. One blunt way to solve this would be to simply keep upping the weight on “test task” till things seemed intuitively satisfying. How else could we solve this?

Normalization and standardization

One way would be to normalize the results. [Here is an example of this on the Hiring Template.](#)

Name	Notes	Offered?	Accepted	Overall additive score	Initial application score	Min-max normalized	Test task	Min-max normalized	Interview	Min-max normalized
				Weights	0	1		5		2
Applicant 1		Yes	Pending response	5.035714286	57	0.75	13	0.8571428571	23	0
Applicant 5		Pending info on logistical issues		6.5	65	1	14	1	26	0.25
Applicant 3					55	0.6875	7	0		
Applicant 2				4.642857143	49	0.5	10	0.4285714286	35	1
Applicant 4					33	0				

In this particular case, I've normalized everything such that it falls between 0 and 1, setting the highest score to 1 and the lowest score to 0 using the equation.

$$x'_{\text{normalized}} = (x'_{\text{raw}} - x_{\text{min}}) / (x_{\text{max}} - x_{\text{min}})$$

such that x'_{raw} is a single datapoint representing a raw score, x_{min} is the lowest score in the dataset, and x_{max} is the highest score in the data set. On Google Sheets, this looks like $= (G3 - \text{MIN}(G\$3:G\$7)) / (\text{MAX}(G\$3:G\$7) - \text{MIN}(G\$3:G\$7))$

After normalizing the score, we see the larger number of points in the “initial application process” no longer distorts the rubric and the results are more in line with common sense.

Another useful method is standardization using a Z-score. [Here is an example on the hiring template.](#)

Overall additive score	Initial application score	Z normalized	Test task	Z normalized	Interview	Z normalized
Weights	0	1		5		2
1.995538198	57	0.4345420758	13	0.632455532	23	-0.800640769
5.205972221	65	1.103068346	14	0.9486832981	26	-0.3202563076
	55	0.2674105082	7	-1.264911064		
0.4266711285	49	-0.2339841947	10	-0.316227766	35	1.120897077
	33	-1.571036736				

Standardization assumes that your data follows a normal distribution, and then represents each data point in terms of how many standard deviations it is from that mean. To standardize your data, use the equation

$$x'_{\text{Z-score}} = (x'_{\text{raw}} - x_{\text{mean}}) / (x_{\text{standard-deviation}})$$

In Google Sheets, you can just use the built-in function

$=\text{STANDARDIZE}(\text{E5}, \text{average}(\text{E\$5:E\$13}), \text{STDEVA}(\text{E\$5:E\$13}))$

Where **E5** represents the individual data point and **E\$5:E\$13** represents all data points in the set.

A Z score assumes normal distribution and expresses every value in terms of standard deviations from the mean							
Normalization sets the minimum value to 0 and the maximum value to 1							
both can be used to combine unlike metrics. Here, I've illustrated that normalization and z scores stay the same despite additions and multiplications to the original value.							
Raw	standardize (z score)	normalize	Raw	Raw * 23 + 102	standardize (z score)	IQ equivalent	normalize
1	-1.686548085	0	1	1044	-1.686548085	74.70177872	0
3	-0.4216370214	0.4	3	1090	-0.4216370214	93.67544468	0.4
4	0.2108185107	0.6	4	1113	0.2108185107	103.1622777	0.6
5	0.8432740427	0.8	5	1136	0.8432740427	112.6491106	0.8
6	1.475729575	1	6	1159	1.475729575	122.1359436	1
4	0.2108185107	0.6	4	1113	0.2108185107	103.1622777	0.6
2	-1.054092553	0.2	2	1067	-1.054092553	84.1886117	0.2
3	-0.4216370214	0.4	3	1090	-0.4216370214	93.67544468	0.4
5	0.8432740427	0.8	5	1136	0.8432740427	112.6491106	0.8
If you find it more intuitive, you can multiply the z-score by 15 and add 100 to convert it into the format of an IQ score.							
Note that real results often don't actually follow normal distributions							

Here is a [spreadsheet](#) you can play around with which normalizes and standardizes a set of values. As you can see, it makes no difference if you arbitrarily multiply the original values by anything or add anything to them – normalization and standardization both remove these sources of noise.

Normalization and standardization can be used together, and don't interfere with each other (i.e., normalizing a dataset with and without standardizing it gives the same answer, and vice versa).

You can use either method, or both methods, as appropriate to the situation. If your data is normally distributed, Z-scores can be a bit more natural and informative, because they represent data in terms of deviations from the mean. If your data is not very normally distributed, then a Z-score would not carry any meaning about standard deviations, making your data potentially difficult to interpret. On the other hand, min-max normalization is very sensitive to the highest and lowest data point – a single extreme value could cause all the other values to be squished closer together. You should generally look at the distribution of your data and try to remove outliers when applying normalization and standardization – or better yet, design your metrics so you would not need to modify your data like this in the first place.

Calibrating multiple ratings

Even after taking care of numerical glitches, you may be concerned that some of your numbers will be subjectively chosen. You can guard against this by having multiple people give ratings.

After the raters have finished, you should highlight any ratings which disagree with each other (if you are using a very fine rating scale, you can highlight ratings which disagree by more than a certain amount) and discuss the reasons for differences in ratings. It helps to rate a small batch of candidates together so you can design a rubric that everybody agrees upon, using the information from the resulting discussion.

These methods can only augment thinking, not replace it

Normalization, standardization, and calibration are not general-purpose strategies, nor are they the only way to deal with the issue of mismatched measurements (for example, you could simply divide overly dominating metrics by some number). The purpose of introducing these methods is not necessarily to encourage the normalizing, standardization, and calibration of everything, but rather to make you aware of how spreadsheet numbers can be misleading. Lack of normalization or standardization is only one of many ways an artifact of calculation can influence your results in a big way. For example, if you calculate scores by multiplying them, you will favor all-rounders over specialists (e.g. $4*4*4=64$ while $1*4*7=28$, even though they both equal 12 when added).

When making a final decision, look over each applicant multiple times, use your common sense, and combine quantitative and qualitative metrics. While you must try to make the numbers correspond in some way to a real-world impact as much as possible, numbers are ultimately only a tool to augment your working memory and decision-making.

Summary

Make use of Google Sheets techniques. You can automatically take a sum of products; avoid having to type equations, with hotkey shortcuts Ctrl-D and Ctrl-R; integrate multiple sheets; and automatically search for specific strings within a set of cells, pulling the data you want from the associated rows/columns (e.g. vlookup). When integrating unlike metrics, use tools such as normalization and standardization to avoid distorting your data. When integrating subjective ratings from multiple graders, use calibration tools to build consistency within your scoring system.

Author's note

I am not trained in statistics. I learned how to calculate a composite score by reading the articles below. If you have a better understanding of statistics and think there is a mistake, please contact ishaan.guptasarma@charityscience.com.

Citations/resources

- [Zapier Guide to Spreadsheets](#)
- [Using Spreadsheets to Make Decisions](#)
- [VLOOKUP like the Pros](#)
- [A two-score composite program for combining standard scores](#). Note: You might notice my equation is slightly different from the one this paper uses. This is because the paper is trying to measure specific cognitive skills, and therefore penalizes metrics that don't correlate with other metrics. (For example, if you wanted to determine English skill by combining performance on two quizzes, you would be happier to trust quizzes that correlate with each other.)
- [Combining scores from different tests](#)
[Combining Multiple Measures into a Single Figure of Merit](#)
[YaRrr! The Pirate's Guide to R – Standardization \(Z – score\)](#)
- [Standard Score – Understanding z-scores and how to use them in calculations](#)
- [Calibration Protocol for Scoring Student Work](#)

5. Weighted Factor Models

Author:

Karolina Sarek,

Co-founder and Director of Research

Charity Entrepreneurship's research process involves **multiple stages** and four methodologies: **cost-effectiveness analysis** (CEA), **expert views** (EpV), **informed consideration** (IC), and **weighted factor models** (WFM). In this article, we explain why and how Charity Entrepreneurship uses WFM as part of our research process, and discuss its strengths and weaknesses as a methodology.

WFM consists of generating a set of criteria with assigned weighting for each and then assessing how a possible option scores on each of these. It is particularly useful as it allows researchers to combine a large number of objective and subjective factors and identify which ones drive the results. However, it has some weaknesses such as its lack of flexibility, which suggests that it is best used in combination with other methods.

CE uses WFM during three stages of our research at different levels of depth: first five minutes, then two hours, and finally twenty hours.

- First, we conduct an **idea sort** to sift out the top thirty or so ideas from our initial list of hundreds. We assess each idea using all four methodologies (i.e., CEA, EpV, WFM, and IC), spending twenty minutes per idea.
- Next we produce a **prioritization report**, spending two hours researching each idea from the idea sort and ordering them according to their promise. We use a different methodology for each cause area: WFM is used for our **animal** ideas. At this stage we only apply one methodology to ensure sufficient depth.
- Finally, our eighty-hour assessments of the top interventions involve all four methodologies.

Concretely, lead researchers use WFM by filling in a prebuilt model taking into account four main criteria: strength of the idea, execution difficulty, limiting factors, and externalities. Each of these will be estimated using included built-in

questions, which are more or less specific depending on the amount of time allocated.

What is a weighted factor model?

Broadly, the process of creating a WFM involves generating preset criteria and weightings and then evaluating how a possible option scores on each of these. WFMs often involve a number of preset criteria ranging from three to twelve factors. They typically generate an endline score based on the option score and the criteria score (normally multiplied together). Both hard factors (such as population size in absolute numbers) and soft factors (such as a score out of ten for population size) can be used in WFMs. The way our team uses WFMs involves pre-generating consistent research questions that are asked across all charity ideas to produce a score for a given criteria.

Example weighted factor model (*Charity Entrepreneurship 2019*)

	Total welfare score (with evidence)	Range	Level of depth	Estimated population size	Odds of feeling pain	Death rate/ reason	Human preference from behind the veil of ignorance	Disease/injury/functional impairment	Thirst/hunger/malnutrition	Anxiety/fear/pain/distress	Environmental challenge	Index of Biological markers of happiness	Behavioral/ interactive restriction	
Max score	100 to -100	100 to -100	Hours	In millions			20	20	17	15	15	5	4	4
Human in a high-income country	81	53 - 90	1	36	99%	17	17	13	13	11	5	2	4	
Wild Chimpanzee	47	13 - 75	1	0.3	85%	6	8	9	11	7	3	0	2	
Human in a low middle-income country	32	23 - 61	1	1324	99%	13	9	6	0	2	1	2	0	
Wild bird	-2	-18 - 45	5	400,000	70%	-10	0	0	4	1	1	0	3	
FF Beef Cow	-20	-58 - 7	3	1000	75%	-2	-7	-4	4	-6	-1	-2	-1	
Wild rat	-28	-48 - 2	3	7000	72%	-14	-10	-10	3	-1	3	0	2	
Wild fish	-31	-51 - -3	5	3500000	60%	-9	-8	-10	0	-6	1	0	2	
FF Cow milk	-34	-65 - -13	5	264	75%	-4	-12	-7	2	-7	-1	-2	-2	
Wild bug	-42	-63 - -4	3	10,000,000,000	10%	-16	-10	-5	-6	-2	-3	0	0	
FF Fish-traditional aquaculture	-44	-58 - -27	3	1000000	60%	-15	-14	-4	4	-7	-3	-2	-3	
EU FF laying hens (enriched cages)	-46	-75 - -31	5	300	70%	-11	-14	-9	4	-9	-2	-2	-2	
Wild fish for human use	-47	-69 - -22	3	0.97-2.7 million	60%	-16	-13	-12	0	-8	1	-2	2	
FF Broiler chicken	-56	-71 - -25	3	22000	70%	-13	-19	-15	8	-9	-4	-2	-2	
FF Turkey	-57	-71 - -26	5	244	70%	-13	-16	-12	1	-10	-3	-1	-3	
USA FF laying hens (battery cages)	-57	-72 - -46	5	260	70%	-13	-17	-12	4	-10	-3	-2	-4	

Example of synthesized expert data (*Charity Entrepreneurship 2018*)

Broad area	Intervention	Ordered Sub-interventions	Weighted average	Average	Range in views
Corporate Outreach	Broad area ranking	NA	2.8	2.7	Small
Research	Broad area ranking	NA	2.6	2.7	Large
Product Creation	Broad area ranking	NA	2.5	2.5	Large
Political Outreach	Broad area ranking	NA	2.2	2.3	Small
Welfare Condition Improvements	Broad area ranking	NA	2.5	2.2	Mid
Wild Animal Suffering (WAS) / Bugs	Broad area ranking	NA	2.3	2.2	Large
Improving Existing NGOs	Broad area ranking	NA	1.8	1.8	Mid
Events	Broad area ranking	NA	1.6	1.7	Small
Veg Outreach	Broad area ranking	NA	1.8	1.6	Small
Other	Broad area ranking	NA	2.0	2.0	Mid

Why is this a helpful methodology?

The WFM is a highly versatile tool because it can incorporate a large number of factors including subjective ones clearly, but simultaneously uses a numerical

calculation to determine the endline result. This allows it to produce surprising results and makes it easier to track down the factors that lead to the result.

Reasons this is a helpful tool (in rough order of strength):

- **Systematism in idea comparison**
- **Enables comparison of all ideas with equal rigor**
- **Reduced gaps**
- **Allows integration of multiple factors**
- **Sandboxing**
- **Allows soft and hard inputs to be combined**
- **More angles for learning**
- **Understandability**
- **Encourages quantified consideration**
- **Can lead to novel conclusions**
- **Makes it easier to communicate conclusions**

Systematism in idea comparison: WFM's encourage considering the same aspects of criteria across multiple ideas. This allows much closer idea comparison than the other models, which each have more idea-to-idea variability. For example, comparing an idea's limiting factor in very similar terms across all charity ideas can lead to a much stronger sense of how well an idea does in this aspect.

Enables comparison of all ideas with equal rigor: We ensure that we apply equal rigor when evaluating the ideas by answering the same research questions that define the criteria in the same way, spending the same amount of time on each idea, and evaluating it in the same way.

Reduced gaps: Many models can be largely affected by unconsidered factors or gaps in the information. For example, if a single important factor was not included in a cost-effectiveness analysis (CEA), it would be hard to detect but could largely affect the results. Due to the same questions being asked across all interventions and the same factors being filled in, there is a lower chance of gaps affecting one idea but not another in a WFM.

Allows integration of multiple factors: Many models are not conducive to including many different factors in a single number. For example, CEAs do not handle **limiting factor concerns** very well unless multiple CEAs are done for many different possible levels of scale. Similarly, many CEAs do not include strength of evidence other than

as a simple discount at the end of the calculation, which does not capture how to weigh different types of uncertainty (e.g., Knightian vs. non-Knightian).

Sandboxing: A large difference between CEAs and WFMs is the total weight that a single factor can hold. In a CEA, one very large number can swamp many small numbers. For example, if an intervention affects a huge number of beings but has a very low chance of working, this initial huge number can make all the other numbers in the CEA trivial. Due to each factor having an effective maximum weight, a single factor affects a WFM far less. You could say the impact of that factor is “sandboxed” within a single factor.

Allows soft and hard inputs to be combined: Some important factors are easy to get a single hard number on, for example “total population affected by measles”; however, other factors are impossible to put a hard number on, for example “the tractability of founding a new charity in India.” These factors can be given a soft number but in a consistent and comparable way in a WFM. These soft numbers can be calculated with harder numbers’ Z-scores to determine what ideas are outliers in terms of many positive factors.

More angles for learning: One of the purposes of our research process overall is to generate better empirical information about how to rule in or out charity ideas more quickly in the future. WFM is the only system we use in which subcomponents could be correlated individually to our endline results. For example, we could determine if the evidence base predicts very strongly what interventions are recommended after deep reports are conducted. Pulling out a single aspect like this from a CEA or expert interviews would not be easy.

Preregistration: In many ways a WFM leaves the fewest areas open to interpretation, with preset questions and descriptions for how different items would score ahead of time. This means that researchers with fairly different starting points and intuitions will more often reach the same conclusions when compared to systems that are more open to researchers’ interpretations. This concern most affects our informed consideration (IC) but can also largely affect CEAs.

Understandability: Intuitive systems can be built into a WFM, making it quick and easy to understand relative to other systems. Color coding is easily used to show areas of comparative strength and weakness across a large number of ideas. Both expert views (EpV) and IC lend themselves to written paragraphs, which are slower. A CEA is quicker to understand the endline number but takes longer than any other system to understand the full logics and weightings behind the numbers.

Encourages quantified consideration: Like CEAs, WFM's encourage quantified and numerical consideration of factors. By default, most people (including experts) do not think in quantitative terms. For example, when asked if an event will happen, most people think of this as a binary question (yes/no) rather than thinking about the probability of the event happening. WFM's require quantitative inputs for each variable, which encourages quantitative thinking and calibration (e.g. an event being 20% vs. 80% likely).

Can lead to novel conclusions: Like CEAs, WFM's can lead to surprising conclusions. With preset methodology and calculations, it is common that after filling out the data, a WFM will suggest something to be high impact that would not have appeared so by taking a softer, more higher-level look.

Makes it easier to communicate conclusions: Because all the factors are researched and scored separately, we can easily distill the advantages and disadvantages of each idea and explain why the given idea is better than the other.

Why it's not our only endline perspective

In 2019 this model was the primary one we used when **comparing charity ideas**, although in some cases we also used **unweighted factor models**. We think that although this model has considerable promise, it also has many weaknesses that can be counteracted by using multiple models. We also see considerable learning value in testing multiple models and seeing which ones best predict our endline mixed model conclusions.

Flaws of WFM (in order of importance)

- **Not a commonly used methodology**
- **Low flexibility**
- **Limited question cross-applicability**
- **Considerable upfront time required**
- **Can make nonnumerical data look numerical**
- **Can be hard to determine source or reasoning of weighting criteria**

Not a commonly used methodology: WFM's are not a commonly used system in the same formal way we use them. Thus there are few established norms and a lower level of initial understanding from both researchers and readers. It also suggests there might be an unknown but good reason why this sort of methodology is not used more often.

Low flexibility: This system is the least flexible and adaptable across different charity ideas with the questions, full methodology, and criteria weightings all preset. This reduces bias but can also give a large amount of weighting to a factor that might be important overall but far less important for a specific idea.

Limited question cross-applicability: A subsidiary concern of flexibility is that specific questions will not be important to cover but research hours will go into them anyway. Likewise an idea that is important but specific to a given charity is less likely to be covered by this methodology.

Considerable upfront time required: A huge amount of upfront methodological time is required when compared to other systems, because most of the methodology is designed ahead of time and closely followed throughout the process. This means that research is not produced for a long time at the start of a research year and also does not yield feedback loops as quickly when updating the methodology.

Can make nonnumerical data look numerical: A concern with the WFM is that it assigns numerical ratings to nonnumerical data. This can confuse and mislead people as to the objectivity of the system if not explained clearly.

Can be hard to determine source or reasoning of weighting criteria: Endline weights are often the only factor closely examined, and due to endline weightings being used to represent a large number of questions and sources of evidence it can be hard to track down what questions factored into this weighting and how heavily each question was factored.

How much weight we give the weighted factor model

Despite its flaws, we view the WFM as a highly important aspect of our process; we see it as having many of the benefits of CEAs but also being somewhat less error prone and likely to have gaps. Ultimately we think that WFMs, as one of our four perspectives, will generally get between one-quarter and one-half of our total endline weighting, with considerable variation depending on the specific charity idea and cause area. We expect WFM to be stronger in areas where there are many different factors at play and limited hard data.

How CE generates weighted factor models

Summary of factors

Related posts are hyperlinked for more information.

- Strength of idea
 - [Cost-effectiveness](#)
 - [Evidence](#) (example [evidence table](#))
- [Limiting factor](#)
 - Funding
 - Talent
 - [Counterfactual replaceability](#)
 - Size of problem
 - Logistical bottlenecks
- [Execution difficulty](#)
 - Difficulty of founding
 - Difficulty of running well
 - Feedback loop
 - Probability of success
- [Externalities](#)
 - Effects within cause area (e.g. [SARP – small animal replacement problem](#))
 - Effects across other cause areas
 - Information value
 - Other

Detailed information on what question is asked to research each criterion and how a score for each criterion is generated are covered below.

Z-score

A Z-score is a numerical measurement, used in statistics, of a value's relationship to the mean of a group of values, measured in terms of standard deviations from the mean. If a Z-score is 0, it indicates that the data point's score is identical to the mean score. A Z-score of 1.0 would indicate a value that is one standard deviation from the mean. Z-scores may be positive or negative, with a positive value indicating the score is above the mean and a negative score indicating it is below the mean.

Z-scores can be used informally to:

1. Standardize values measured across multiple different criteria, so they can be combined into an overall score and compared to other ideas. For example, we can have an overall Z-score for a given idea based on how it compares to an average in terms of CEA, expressed in \$ per DALY; population size affected, expressed in millions; and crowdedness, expressed in percentage of the problem addressed by other entities.
2. Assess how a given idea scores compared to all the other ideas considered (including an average idea), for example, idea x is better than 70 percent of the ideas on our list.
3. Spot anomalous values. For example, if one of the factors in the scale were an objective number such as population size, a Z-score value would show which countries are outliers relative to others even though population size can differ by orders of magnitude.
4. Reduce risk of some biases, for example, in a situation where the score is not converted to a Z-score, we may happen to use a higher range of values for one criteria but not for another one, effectively changing its weight. For example, suppose a given intervention is evaluated on each factor on an arbitrary scale of 1 to 10. However, one criterion, scale, varies significantly, and you tend to give out sevens and eights frequently, while on the criterion of tractability, you tend to give very consistent scores of four or five. The net effect is that even if you think tractability is more important, you end up weighting scale higher. The method of converting this to a Z-score takes care of this.

More on Z-scores can be [seen here](#) and in “[The Failure of Risk Management](#)” by [Douglas W. Hubbard](#).

Color coding

Color coding is used throughout the spreadsheet to increase ease of reading, with red values generally being areas of weakness and green values being areas of strength. This can allow the reader to quickly see which areas to look deeper into and which areas lead to the resulting total score.

	Program	Overall	Flexibility	Cost effectiveness	Metric focus	Logistical possibility	Indirect effects	Evidence	Scalability	Other factors
Malnutrition	Iron supplementation	8	Mid	Mid	High	Easy	High	High	High	Mid
Cash	Conditional cash transfers	8	High	Low-High	Low-High	Mid	High	High	High	High
Malnutrition	Super salt charity/multimicronutrient	8	Mid	High	Mid	Easy-mid	High	Mid	High	High
Behavior change	Mass media to promote behavior change	8	High	Low-High	Low-High	Easy-mid	High	Mid-low	High	High
Immunization	Immunizations	7	Mid-low	High	Mid-low	Hard	Mid	High	High	Mid
NCDs	NCD best buy	7	High	High	Low	Mid-Hard	High	Mid-low	High	Mid
Malnutrition	Treatment for severe acute malnutrition	7	Mid	Mid	Low	Easy	High	Mid-low	High	Mid
Malnutrition	Vitamin A supplementation	7	Mid	High	Mid-low	Easy	Low	High-mid	Low	Low
Data	Charities that collect or generate information and data relevant to GW	6	High	Unknown	Mid	Mid	Low	Low	Low	High
Malnutrition	Salt iodization	6	Mid	High	Lowest	Easy	High	High	Mid	Mid
Behavior change	Breastfeeding promotion	5.5	Mid	High	Low	Mid-Hard	Low	Mid	Mid-High	High
Malnutrition	Folic acid	5.5	Mid	Mid-low	High	Easy	Mid-low	High	Mid	Mid
Behavior change	Handwashing promotion	5	Mid	Mid-high	Mid-low	Hard	High	Mid-low	High	High
Targeted treatment	Pneumonia treatment	5	Mid-High	Mid	Low	Mid-Hard	High	Low	High	High
Diarrhea	Oral rehydration therapy	5	Mid-low	Mid-low	Low	Mid-Hard	Mid	High	High	Mid
Mass treatment	Deworming	4.25	Mid	High	Mid	Easy	High	Low	Mid-low	Mid
HIV	Provision of anti-retroviral therapy to treat HIV/AIDS	4	Low	Low	Low	Hard	High	Mid	High	Mid
Diarrhea	Therapeutic zinc supplementation	4	Mid-low	Mid-low	Low	Mid-Hard	High	Mid-low	High	Mid
HIV	Prevention of mother-to-child transmission of HIV	3	Mid-low	High	Low	Mid	Mid	Mid	Low	Mid
Targeted treatment	Trachoma control (mass drug blind)	3	Mid-low	High	High	Hard	Mid	High	Mid-low	Mid
Targeted treatment	Tuberculosis case finding and first-line treatment	3	Mid-low	Mid	Mid-low	Hardest	Mid	High	Low	Mid
Cash	Unconditional cash transfers (external)	3	Mid	Low	High	Easy	Mid	High	Mid	Mid
Malaria	Malaria treatment	3	Low	Mid-low	Low	Hard	Low	Mid	Mid	Low
Mass treatment	Mass administration of ivermectin and albendazole to control lymphatic filariasis	2	Mid	Lowest	High	Mid	High	Low	Mid-high	Mid
Immunization	Immunization to prevent maternal and neonatal tetanus	2	Mid-low	High	Low	Mid-Hard	Low	High	Low	Low
Mass treatment	Mass administration of ivermectin to control onchocerciasis (river blindness)	2	Low	High	Mid	Mid-Hard	Low	Mid-low	Mid	Mid
Mass treatment	Polio eradication	1	Low	Unknown	Mid	Hard	Low	Mid	Very low	Low
Malaria	Bednet distribution	1	Low	High	Low	Easy	Mid	High	Mid	Low

Different lengths of weighted factor model

Five minutes

Process

In five minutes, the larger-scale questions cannot be considered in depth. Instead, each factor can be intuitively considered and given a ranking; these rankings can then be added together, resulting in a total score. The factors should be understood first (by reading the content in this document and the linked content for each of the metrics).

Questions to consider

- How well does the intervention do on only this factor?
- Do I expect this intervention would score well or poorly on this factor upon deeper research?
- Does thinking through each of these subfactors change the score?

Expected outcomes

- Four subjective ratings (one for each of the large categories)

Advice

- Consider looking at an important factor as a whole for a few hours and then attributing the ratings to all the interventions for that factor. For example, look at funding allocation in the space for three hours, and then rate all the ideas on availability of funding.

- This is less relevant for a factor such as strength of the evidence, for example, for which you will be better off going through each intervention one after the other and spending one to two minutes looking for supporting studies on Google.

Two hours

Process


At the two-hour stage, each question can be considered from the full **set of prebuilt questions** with deeper research occurring for the most important question within each section. In this stage the questions work as more of a guide than a necessary list to answer. (For our 2020 research, this stage is only applied to animal advocacy interventions. Other cause areas employ a different two-hour methodology.)

Questions to consider

The weighted factor model questions

Expected outcomes

- The first-page summary filled out, connected to sources and ready for polishing
- A description for each section and subsection and why it was scored that way
- A simple causal chain (we suggest using **draw.io**), for example:



Twenty hours

Process

At twenty hours, each question in the document should be considered, researched, and answered. Ratings should be given thoughtfully and updated as new evidence comes in.

Template for 20h WFM:

Three-sentence description of how the intervention does in each criteria area.

Criterion	Ranking	Description
Strength of the idea	/10	
Limiting factor	/10	
Execution difficulty	/10	
Externalities	/10	

Basic causal chain:

These questions are guide questions and should be answered, even if tentatively. More questions can be added and answered that seem applicable for the specific charity idea.

Criteria

Strength of idea questions

Key question: When you look at the theory of change that includes assessment of the strength of evidence, and rough cost effectiveness model, how promising does this idea look?

Theory of change and cost-effectiveness

We suggest implementing this section at the end of the factored model. NB do not make a cost-effectiveness model because this will be an entirely different section. This section is more to flesh out other people's CEA models and plausible theories of change that will then be used to create a CEA.

Key question: What is the plausible path to impact of this charity idea? What do current estimates or expert views of cost effectiveness look like? Does it seem like it compares favorably to other ideas in the area?

- Consider possible paths to impact
 - How are other actors in the space making an impact in this area? If you made three different possible paths to impact, what would each look like?
- What is the causal chain of this charity idea leading to impact in the world?

- Create a theory of change including long-term effects, and list the evidence base for each step including soft steps like **chance of sentience** (**best resources**). If no empirical evidence can be found for long-term effects, note how confident you are in your predictions about it.
 - What aspects of it are least certain? What aspects are most important? Research each of these.
- Are there any alternative ways to approach it that have the potential of making it more cost effective? Such as...
 - Only doing a part of the intervention (e.g., marketing not inventing new in vitro meat)
 - Partnering with another organization
 - Pairing it with other interventions that have similar distribution channels
 - Implementing it on a mass scale, even if it is only effective for a percentage of the population
 - Starting in the country that could absorb change faster
- What metric will make the most sense to use when making a formal CEA?
 - Welfare points? DALYs? SWB?
- What do the experts think about the cost effectiveness of this intervention? Are there any CEA models created by other organizations or costs that can be pulled out of studies?

Evidence

Key question: Overall, how well-evidenced does this intervention look? Does the supporting evidence come from many different robust sources?

- Experts
 - What do the experts think from reading their online content?
 - What specific pieces of data do experts point to?
 - Search for critiques of this intervention. What do people say is wrong with it?
- Quality and quantity of studies
 - Approximately how many randomized controlled trials (RCTs) or other well-designed studies have there been on this intervention or related interventions?
 - Format them into an evidence table **like this one**
 - Use this **template** for assessment of evidence base
- Historical evidence
 - How many charitable dollars have been spent on this intervention?

- Has there been historical success of interventions like this one?
- How much good has been done with the average dollar donated to this cause historically?
- Does this intervention seem responsible for a **disproportionate** amount of impact per spending put into it?
- Are there any situation-related factors that should be discounted?
- Other sources of data
 - How direct is the effect?
 - Are there any positive macro-level data?
 - Is there any empirical analysis of broad (e.g., country-level) trends that could inform us about the effectiveness of this approach?
 - What would we see in a world where this sort of intervention works? What would we see in a world where it does not?
 - Do these macro-level data support what the experts claim about the impact of historical work?

Evidence: Robustness

- Does the intervention look good across multiple values and epistemic views?
 - Moral views: e.g. **classical utilitarianism**, **prioritarianism**, **negative utilitarianism**, **philosophical majoritarianism**
 - Epistemic views: e.g. empirical, skeptical, rational, fast acting, cluelessness
- Does this intervention rely heavily on one type or piece of evidence?
- Could many things change in the evidence of this intervention that would change the endline conclusion on it?
- Would this intervention still do well on other plausible metrics outside of the one selected for these questions?
- Are there any other sources of confirming or disconfirming evidence that could be researched? Brainstorm at least five and research the top two.
- Do multiple pieces of evidence converge on the same answer or do they point in opposite directions?

Limiting factor questions

Key question: What is the main limiting factor to scaling this intervention? At what size does it cap the intervention?

Funding availability

- Is there reason to think that this intervention will be difficult to market?
- How much will this appeal to our different group of funders? Within EA? Tangential to EA? Disconnected from EA?
- What is the total amount of funding in this area?
- How much would it cost to run this intervention at the smallest level of scale. Is funding available for that?
- What is the probability that a new funding pool would appear after introducing an intervention?
- How high is the bar for funding? Are weak projects getting funded in a similar space?
- How hard or easy would it be to build a funding base up over time?

Talent availability

- How specific and difficult-to-find a skill set does this charity need in its staff?
- What skill sets are required to execute this intervention successfully? Are the skills overlapping with the talent gaps in the EA movement?
- How many existing charities working in the same field are hiring for leadership and management positions?
- How many existing charities working in the same field are hiring for positions that are crucial for the success of this intervention?
- Could the talent gap be closed by non-EA/AR experts?

Counterfactual replaceability

- Are there any strong charities working in this area?
- Are there many other people/groups in the area?
 - Who are the other major bodies working on this intervention?
 - How many standard charities seem to be working in the area?
 - Are charities activity planning on moving into the area?
 - How common is this specific intervention in nonprofits that work in the cause area?
 - Is there a government program in place doing this intervention?
 - Are new orgs getting founded in the area?
 - What percentage of this problem seems to be uncovered by other actors in the space (a rough idea)?
- Are they big and competent?
 - Are they doing a good job?
 - Are they doing what seems to be the most cost-effective intervention?

- Do they seem scientifically minded and aiming to do the most good?
 - What is your rough impression of them?
 - What areas are covered? What do the coverage rates look like? Is it money or is there some other obstacle to full coverage?
- Will they generate relevant research that your group can use as time goes on?
- Could organizations/groups be influenced to do a better job?
- If the area is not crowded, why not?
- Is attention paid to this invention growing? How fast?

Size of problem

- On what scale could this intervention be delivered? Is it an ultimate limit or just a matter of diminishing returns?
- Approximately how many animals/humans are affected and what percentage of the problem is covered by other orgs?

Logistical bottlenecks

- What is the total number of hires who would be effective in this area?
- Is the intervention a one-off thing or does it need to be regularly repeated?
- Would this need to be evaluated with an RCT before scaling?
- How fast could this be scaled?
- Could it be run not only by employees but also by interns and volunteers?
- Are some parts of the intervention easy to automatize?
- Would it be easy to build a community around the organization?

Execution difficulty

Key question: Overall, how hard is it to set up and run this intervention well relative to others on the list?

Difficulty of founding

- Would a charity in this area have to establish any major partnerships with corporate or government bodies, or could it be mostly run with in-house staff?
- What do experts see as the biggest challenge of running this intervention?
- Cost structure:
 - Where will the money be spent on this intervention? What is the rough breakdown of costs between staff, materials, manufacturing, logistics (supply chain and distribution), technology, and administration?

- How much of the work of administering the program can be performed by employees recruited locally?
- Are any of the steps provided for free or at discounted rates? (e.g., flyers can be provided for free by other NGOs; Google grants)
- Can we bring down costs over time with scale, technology, or the right expertise?
- Will this area seem intimidating to founders?

Difficulty of running well

- Do experts consider executing this intervention to be easy or difficult relative to other interventions?
- How many people would the charity have to interact well with?
- Are there any cultural factors that may make performing this intervention difficult?
- What is the biggest barrier to successful implementation?
- What have been the stumbling blocks for charities/governments who have tried this intervention before?
- Are there any unexpected factors that make this intervention vulnerable to lower impact if things don't go as planned? For example, reliance on tech, need for highly skilled people, low uptake of intervention, etc.
- Does it have value drift, affecting factors that might make a strong EA founder lose their focus on altruism?
- Is cost effectiveness highly sensitive to details of ongoing staff decision-making. i.e., will it fail to be cost effective if they make a subtle wrong call on a, b, c, d, e, f...?
- Is there anyone (private persons/farmer's associations/political bodies/other activists) who would be opposed to this intervention? How much influence and power do they have? Do you see a potential for conflict resolution?
- Have past projects in the area failed?

Feedback loops

- How hard is it to test the effectiveness of this intervention?
 - How feasible is it to run micro pilots or RCTs on this intervention?
 - How quickly could this intervention be tested? (e.g., how long does treatment take?)
- How clear are progress metrics on a short time frame, e.g., month to month?
- What would be the evidence that would show this charity is working
 - Well?
 - OK?
 - Poorly?

- How easy would it be for an outsider to see whether the charity is performing well?

Probability of success

- Is this intervention hits-based or reliable in generating results?
- What has the historical hit rate been for other groups doing this?
- Is the expected value of this a power law distribution or bell curve distribution?
- Is the success metric more robust or more fragile? Does the metric rely on many assumptions to count as a “hit”?

Externalities

Key question: What other possible negative and positive effects will this charity have? How large are they estimated to be? How much evidence is there? How much confidence do you have in the effects?

Within cause area

- Could this intervention turn people off the cause area as a whole?
- Could this intervention affect group interactions within the cause area in a negative way?
- Look for data on whether anyone has reported being negatively affected by this intervention in the past for any reason.
- What is the risk of this intervention causing damage within the cause area?
- Can this risk be mitigated? If so, how?
- What is the risk of value drift for co-founders who work in this area?
- Relative to other ideas in the area, will it be easy or hard to be transparent while implementing this idea?
- Could a field be largely changed in a more positive direction if a charity was founded in the area?
- What are the odds this field will grow significantly in the next ten years? Fifty years?
- Does this intervention have a net effect of support or harm for other charities in the area?
- If successful, would the actions increase the probability of success in similar actions of other groups? (e.g., legislative change in one country would make it easier to change the law in another country.)
- Would starting a charity in this area be generally recognized as a good thing? From supporters of this cause area?

Outside of cause area

- Are there any huge flow-through effects that could dwarf the direct effects of this cause (e.g., **small animal replacement problem**)?
- Are there any significant positive effects that could improve the impact (e.g., environmental impacts, animal impacts)?
- Does it affect the culture's morals in a positive direction or expand the moral circle?
- Does this intervention have effects in terms of helping humans' GDP or speeding up science and technology?
- Does this intervention promote cross-applicable ideas such as anti-speciesism that could help in the far future?
- Is there any evidence that this program will reduce wild animal suffering (including bugs)?

Information value

- Relative to other interventions based on priority programs, does this intervention provide the charity world with learning value? (e.g., additional evidence)
- Is the field young or new?
- Would starting a charity in this area give you skills you could pass on to other charity entrepreneurs?

Other

- Will this charity be able to help other new charities in the future?
- Is now a good time to start this charity vs. starting it at any other point in the future?

How WFM compares to other methods used in an 80-hour report:

- [10 hours] Broad undirected reading and **crucial considerations** (informed consideration)
- [16 hours] Directed research (weighted factor model)
- [10 hours] Finding and talking to experts (experts)
- [20 hours] Cost-effectiveness analysis creation (CEA)
- [4 hours] Directed research (weighted factor model)
- [10 hours] Summary writing and internal contemplation (informed consideration)
- [10 hours] Showing endline report to experts (experts)

Expected outcomes

- The full document of questions filled out in sufficiently clear detail that a reader could follow it
- A description for each section and subsection and why it was scored that way
- A strong theory of change model
- Filled-out details in the charity idea comparison spreadsheet

Reading list

Internal resources

Our process for narrowing down which charity ideas to research

Metrics

Cost-effectiveness

The Importance of Evidence

The importance of being flexible

Why you should care about scalability

Why you should care about indirect effects

How logistics will influence which intervention you pick

Counterfactual Impact: What Would Happen If You Didn't Act?

Why We Look at the Limiting Factor Instead of Problem Scale

Using a Spreadsheet to Make Good Decisions

External resources

Sequence thinking vs cluster thinking

Larrick, Richard P. "Broaden the decision frame to make effective decisions."

Handbook of Principles of Organizational Behavior (2009): 461–480.

III. Tools of rationality (the hammer)

The humble hammer is a tool that most are familiar with, many will have used, and few are experts with. A hammer can be used to quickly make progress on something that would be very hard to do without, and hammers are a key tool in building new things. However, a hammer can also make someone overconfident. Many people think they are an expert builder if they own a hammer, and hammers can cause a lot of damage if used improperly.

Rationality is a powerful instrument that can be learned to improve a huge number of decisions. Techniques can often be taught in a single blog post and can rapidly lead to better choices and to avoiding obvious mistakes. However, knowledge of rationality runs far deeper than we are able to go in this book. Often due to its intuitive importance it is overused in situations when it should not be. In this section, we will cover:

- The idea of skepticism applied to your charity
- Applying rational thinking to calculating how much time you have to make an impact
- How rationality can be used to find and work on the crucial considerations that matter most in an issue.

1. What Does It Take to Be A Skeptical Charity?

Author:

Joey Savoie,

Co-founder & Director of Strategy

Organizational self-confidence is seen as a very positive trait and is often embraced by new start-ups and nonprofits. The concept of **self-skepticism** is arguably as important but is often neglected. Self-skepticism cautions us to clearly measure our impact before declaring ourselves effective or expanding our organization. We believe that a healthy dose of this applied to the charity sector would make the world a much better place much faster.

Below is an outline of our organizational self-skepticism checklist.

It breaks down into two main categories: (1) **do not assume a priori that the organization is better than the evidence suggests**, and (2) **constantly seek and use critical feedback**.

An organization does not assume that they are better than the evidence suggests, by...

1. Not assuming they are far outside the **base rate** or norm without strong evidence.
2. Not assuming that the average and the marginal organizational impact will be the same.
3. Only seeking major funding/expansion/scaling once there is evidence of impact.
4. Individuals within the organization being skeptical about personal impact until proven otherwise.
5. Being very careful not to generalize strength in one area to strength in many/all areas.
6. Not assuming they are having a large impact in all unquantified areas.

An organization seeks and uses critical feedback, by...

1. Being **transparent** and seeking feedback on their work.
2. Being clear with the public about their goals, objectives, and plans; and seeking feedback on them.
3. Seeking public feedback from critically-minded others about their organization and core strategy, and sharing this information publicly.
4. Getting external reviews and impact evaluations written by critically-minded people not working with or for the organization. Applying major changes when negative results come back from external or internal reviews.
5. Encouraging feedback on organization and addressing it instead of ignoring or squashing it.
6. Keeping open the possibility of shutting down if **impact is not proven**.

We feel as though **one of the largest realizations that has made us far more effective was to start applying a much stronger self-skepticism to our own ideas**. Often, we can be harsh and critical when it comes to others' ideas, but have a much harder time seriously evaluating our own impact.

2. What If You're Working On the Wrong Cause?

Author:

Kat Woods,
CE Co-founder

If your main contribution to EA is time, how long should you spend trying to figure out the best thing to do before you switch to taking action? The EA community has spent a lot of time thinking about this question as it relates to money, but money and time differ in important ways. You can save money then give it at a later date; you cannot do the same with time. In this article I will show my current best guess at the answer to the question.

Broadly speaking, you should switch from researching to acting once the expected value of marginal research equals the expected value of acting. The goal then is to figure out the values of these two parameters.

The value of taking action depends on:

- Time left. The expected amount of time you'll be able to capitalize on your research. This is affected by things like:
 - You changing (i.e. value drift)
 - You retiring
 - The world changing
- Inspiring others. How many other people do you inspire to act based on the example you set?

The value of marginal research depends on:

- Value increase. How much better does your best option become after researching the possibilities?
- Delegability. If others want to enact your conclusions as well as or better than you would, you should research indefinitely, as long as you're still making progress.
- Research progress. What's the shape of progress? Is it an **S-curve**, with tons of low-hanging fruit at the beginning, then getting progressively harder to find subsequent discoveries? If so, what's the slope of the curve during its most productive interval?

The conclusion I drew from these considerations was to invest heavily in up-front research, then do research at spaced intervals to account for considerations you missed, new ones others have thought of, and the world changing over time.

The initial up-front research time can be calculated by putting up the above considerations into a formula based on your best estimates, then figuring out where the marginal value of research dips below enacting the conclusions you came to. Our current best guess suggests we should spend two to eight **person-years** researching.

What's the existing literature on the topic?

There is **a lot already written on doing good now or doing good later** in the EA community, but mostly in regards to giving money. There is also much written about the topic in the wider decision theory community, where it's commonly referred to as optimal stopping or explore/exploit trade-offs. While there are many interesting ideas in the area, their solutions cannot be straightforwardly applied to EA because they solve problems fundamentally different from those we face.

The secretary problem is probably the most famous example of an optimal stopping problem, but it is not a good fit for analyzing EA decisions. Briefly, the thought experiment sets out to figure out how many secretaries to interview before you hire one. Given the conditions of the scenario, there is a mathematical solution – interview 37% of the potential secretaries during the time you can afford to spend interviewing, and after the 37% mark, hire the first secretary who is as good or beats the best found in that exploration phase.

The reason that this cannot be applied to our question is that it assumes that you can quickly tell which secretary is better than another, but in EA, problems are very difficult to compare. For example, is deworming or bed nets better? It's very unclear, and that's in a relatively well studied area. Comparing animal rights research to preventing AI x-risk is even more fraught with ambiguity. Other limitations of the secretary problem are discussed in the comments of **this post here**.

The multi-armed bandit problem is similarly limited. If you are at a casino with multiple slot machines (sometimes called one-armed bandits), and you don't know each machine's probability of payoff, which arms do you pull and in what order? There are multiple solutions to this problem and its variants, however its applications to EA are limited.

For example, it assumes that when you pull an arm, you instantly know what your rewards are. This is a hard assumption to make in EA. Even if we had perfect knowledge about the results of different actions, it would still be unclear how to value those results. Even if you are an ethical anti-realist, there could be considerations that you hadn't thought of before that change the size or even the sign of your expected effect (e.g., your stance on speciesism or population ethics). Despite these and other unlisted problems, there are still some useful takeaways from the literature. Those I found most useful were the arm analogy, the Gittins index, and the observation that explore/exploit decisions depend largely on the time you have left.

What counts as “pulling an arm” in the multi-armed bandit?

In the multi-armed bandit problem, the **Gittins index** is a solution. It says to pull arms you haven't pulled before, because they could reveal an arm that will beat your current best. However, if after you've pulled the arm a certain number of times and it still hasn't beat your current best guess, you can move on.

So what counts as pulling an arm when it comes to EA? Is it starting a charity or working for an organization? This doesn't seem right. Even though I have never worked for **Homeopaths Without Borders**, I can know that utility payoff won't be good. I already have, in some sense, pulled the arm.

This leads to the idea that pulling an arm is analogous to activities that gather information about the option under consideration. Thus the most natural equivalent to pulling an arm is doing a unit of time of learning. The unit of time is relatively arbitrary and can be cut up into very large or small amounts, so I'll just use a day for simplicity's sake. Learning can be done through doing the option or researching through other methods.

This speaks to the question of how long to give a cause a chance before giving up. For example, if you are not convinced of a cause in the EA community, how long should you keep reading new articles about it, engaging in debates with its supporters and so forth, before you stop and start researching other possible contenders? A concrete example would be, if you know a lot about NTDs, but know very little about international trade reform, it would be more valuable to research the latter more.

Explore/exploit decisions depend on the time you have left

Another useful heuristic that the optimal stopping literature gave is that how long you spend exploring versus exploiting depends on how long you have left to exploit your best option. This fits with intuition fairly well. For example, if you are on your deathbed, you probably shouldn't waste any time in trying to make new friends, but rather use your precious last hours with people who you already know and love. However, if you're in college and have many decades left of your life, you should probably invest a lot of time making friends, because if you find a great new friend, you will be able to enjoy that relationship for decades to come.

This applies to charity in the sense of how many working years you have left, although to be more specific, it's not how many working years you have left, but how many years you'll be able to capitalize on your research.

Value of doing

Time left

After brainstorming, we came up with 10 factors that affect how many years you should expect to be able to capitalize on your exploration phase. I expect these to vary enormously based on personality, history, choices, environment, etc., so one person's answers cannot be generalized to other people. Nonetheless, you can definitely apply the same process to yourself as we did and see what the results are. The factors are:

1. **Expected retirement age.** When do you think you will retire?
2. **Health degeneration.** This is a minor factor in the developed world and affected our calculations negligibly. However, this could be different based on your own risk factors.
3. **World change.** The world probably will change after you've done your research and switched to action. What odds do you place on the changes affecting your action to the point where you have to go back to the drawing board?
4. **Value drift.** How likely do you think that you will stop being an EA due to changing your mind, losing motivation, etc.?
5. **Research obsolescence.** What odds do you put on finding a new body of ideas/research that completely obsoletes your previous research? For example, many of my altruistic plans and research prior to EA got completely trashed upon learning about concepts such as counterfactuals and the base rate fallacy.

6. **Pinker Effect.** Pinker's book, *Better Angels of Our Nature*, makes a great case that the world is getting better. Maybe if we spend too much time researching, the remaining interventions will be less effective than the ones we know about now.
7. **Inspiring others.** You may be able to inspire people who will roughly, or very closely, enact your values, thus increasing the amount of person-hours that your research could capitalize on.
8. **Greater hiring ability.** If you gain an ability to hire people, you can further increase your ability to affect the world.
9. **Close-mindedness freeze.**⁴ It is a common conception that the elderly are more close-minded than the young. If you think your mind will effectively close early on, you should make sure to do all of your research before that, so you don't lock in on the suboptimal conclusion.
10. **Flake drift.** This is the counterpart to close-mindedness freeze. It's the idea that even if you come to a great conclusion, maybe personality-wise you can only stay on any given project for a certain amount of time before you get bored and move on.
11. **Unknown unknowns.** No matter how thoroughly you've thought through something, there are always things you haven't thought of.

I will cover some of these factors up close.

Value Drift

Theoretical approach

The most commonly cited concern about how long you will be able to capitalize on your research is value drift. Many things could cause this such as burnout, having children who then take priority, getting distracted, etc.

An important thing to keep in mind with value drift is that there are degrees of it. Scaling back your involvement by 10% is not the same as giving up on EA altogether and becoming a surf bum in Hawaii. Burning out so that you need a two week vacation is different from so thoroughly burning out that you never give another hour of your time.

The risk of value drift is a very personal one so cannot be generalized easily, but by the same token, it is very easy to be biased towards oneself. People typically have

⁴ While this isn't technically about how long you can exploit your best choice, I thought it was relevant so included it anyway.

the end of history fallacy in terms of their personality, **consistently under-predicting how much they will change in the future**. In fact, since I've joined the EA movement I've seen a substantial percentage of people be very enthusiastic and involved at the beginning, only to completely switch or lose motivation a few months or years later. The movement is still very young, so I suspect an even larger proportion to leave as time goes on.

Likelihood of value drift is influenced by:

- **Community.** People are very influenced by their peers. The greater a percentage of people you hang out with who share your values, the less likely they are to shift.
- **Career capital.** If you build up your career capital that's only really useful in the charity field, you're less likely to drift because switching sectors would set you back to square one.
- **How long you've been an EA.** If you meet somebody and they say they just decided to go vegan that very day, what odds do you put on them being vegan 10 years from now? You probably put them very low, and that makes sense. If they've already been vegan for 30 years, you'd put them much higher. Likewise with EA. If you've been an EA for 10 years, you'll likely stay that way. If you're new, you should probably put high odds on your values drifting, even if you are really excited at this very moment.

Empirical approach

Are there any empirical studies that shed light on the issue of value drift?

Unfortunately, there is little data on the issue. There were some interesting studies on how many people who became social workers stayed in the field, but the literature was inconsistent and the measurement only a rough proxy. For example, if somebody leaves government work to run a nonprofit women's shelter, does that count as leaving social work? Likewise, what's the relevant reference class for EAs leaving the movement? Should I put myself in the category of those who donated \$20 to AMF then forgot about GiveWell? Or maybe it should be the people who've started charities in the area? That seems like reference class tennis to me, and I do not have a current solution to that issue. In the end, the empirical approach did not provide much new information.

At the end considering all of these factors, for myself, I put a 15% chance on major value drift, and 45% chance on minor, with both of these most likely to happen earlier on and less likely as time goes on. This had a large effect on my endline predicted working years left, which is to be expected.

Pinker Effect

The world is getting better, which is fantastic, but could it be bad for our altruistic endeavors? Could we run out of all of the good options? I think the answer is no, unless you're working exclusively in global health.

Let's start with global poverty, especially global health. It is undeniably getting better and at an incredibly fast rate. It would be unsurprising that 60 years from now, it will be considered strange that anybody went without bed nets or their recommended vaccines. If at the end of your research, you decide that global health is your top cause, you should get started right away, as all of the good opportunities are indeed getting snatched up.

However, global health is not the only cause. Some causes are not getting better but worse, **such as animal welfare** or environmental degradation, so there might be more to be done to help in the future.

An additional benefit of the world getting better is that we're getting better at helping too. There might actually be more effective interventions in the future because people have spent a longer time thinking about and testing different strategies. For example, medieval activists didn't have the opportunity to provide vaccines because they didn't exist yet.

In the end, the biggest factor for us is that we are becoming wiser and expanding our **moral circle**. There are likely other causes, such as bug suffering, that could be extremely valuable and neglected, that society will ignore, like factory farming, for decades or centuries to come.

So, the good news is that the world will still have problems after your research*, so don't worry too much about it getting better. You should probably not have this dramatically affect how many working years you have left.

*Just as a note, in case it wasn't clear because tone can be lost in writing, I am 100% joking that it's good news that there will still be problems when we're older. It would obviously be great news if there were no more suffering.

Potential for a larger team/inspiring others

I have the fortune to be on a team of like-minded individuals, such that we can have a high level of coordination. This means that I can focus on research while others

focus on doing the action that we currently think is the highest value. The higher alignment this is, the more my effective sphere of influence is. If I can completely rely on one other person, and them on me, we can get twice as much done as a single person.

This is true at the very high level of alignment with a small number of people, but could also be true with a large number of people but with less value and epistemic alignment. To illustrate the point, imagine you do the research and it inspires somebody to start the exact charity that you would have wanted. Action is highly delegable, and if your comparative advantage is research, you should continue researching and then propagating your research, advocating for others to act upon it. This is the general strategy of GiveWell when it comes to recommending where others give, and 80,000 Hours in terms of its career research.

The question is – how delegable is action? Money is fairly straightforward. A dollar given to AMF by somebody who hates science does the same amount of good as a dollar given by a science geek. However, if a charity is run by one person instead of the other, many different choices will be made that will affect the effectiveness of the charity. For example, the science-disregarding person might hear anecdotes of bed nets being used as fishing nets and switch to a different intervention, whereas the science geek might read the literature and see that while this occasionally happens, it's swamped by the positive effects.

There are some examples in the EA community of researchers inspiring charities, and founders stepping back from their roles as CEOs, which can provide a rough outside view. You can check how valuable the charity continued to be, according to the founder's values and epistemics, compared to how it was or would have been had they continued to run the charity. Based on the examples I am familiar with, approximately 15% got better after handing off, 45% stayed the same, and 40% got worse.⁵ However, this changes if you take into account how much time the founder or researcher invested in the charity at the beginning, ranging from simply writing a blog post about the idea to spending multiple years setting up the organization. With heavy investment, 0% got worse, 85% stayed the same, and 15% got better. There are other factors aside from founder time that affect delegability:

- **How good your ideas are.** If your ideas are terrible to everybody but yourself, you will have to enact them yourself because you won't be able to persuade anybody to follow your advice.

⁵ For obvious reasons, I cannot publicly elaborate on what these numbers are based on.

- **How persuasive you are.** Even if your ideas are great, if you find it difficult to persuade people of simple things, persuading them to start a charity based on your advice will be impossible.
- **How palatable your ideas are.** If your best option is starting a global poverty charity, most people think that is a worthy cause. If your best option is fighting factory farming, many people think this is not actually a problem, so there will be fewer people willing to organize against it.

Despite our relatively pessimistic views on delegability, this still represented a huge increase in our number of “effective working years” and thus the value of upfront research.

Conclusion on working years remaining

To put together all of these considerations, I started off by assuming that I would retire at the normal age, due to factors pulling for and against late retirement. This left me with 40 years. Then I took off or added expected years of work, based on probabilities I put on the different factors. Results will vary based on your personality, choices, and environment. For myself, after putting in hours and hours of work and thought and calculations, I ended up, anti-climatically, with 40.2 years of expected work.

This was not what I was expecting, but it was still worth the effort. Initially I had simply thought about value drift and applied a steep discount to my work, but I had not taken into account any positives, or thought about the whole picture. I recommend others trying this exercise as well, because it could affect your decisions.

Flow-through effects

It has been argued that **the effects of doing good now compound**. That if you inspire one person to do earning-to-give, then you will continue inspiring new people, and they will too, thus “earning interest on your interest”. I believe this is an oversimplification of the effects.

For one, say you start a direct poverty charity and that inspires approximately one new charity per year, and those charities have the same “inspiration rate”. This won’t go on forever until everybody in the world is starting direct poverty charities. It’s not an exponential curve, but rather an s-shaped curve. There are the initial low-hanging fruit, exponential growth for a period of time, then a diminishing amount. However, this isn’t the end of the story. After all of these charities start,

they don't last forever. People retire, charities shut down, problems are solved, etc. So really after the tapering, there is a probably relatively linear comedown.

Additionally, compounding benefits apply to doing good later as well. It's not like if you start a charity 10 years from now, nobody will care anymore. However, there is still a penalty for starting later. If you spent 39 years researching, then spent 1 year doing, you'd only have 1 year of inspiring, so only one extra charity started because of you. Compare that to if you had spent 1 year researching and 39 years doing, at which point you'd have far more charities inspired by you.

It is important to note that this model compares acting now, or doing the same action except years later. This means it does not take into account the increased value of your best option that you reap from more research.

Furthermore, this is probably an overly optimistic scenario. There are many more ways to deviate from doing good through doing rather than giving. If you are earning to give and you inspire another person to earn to give, if they donate to the same charity or set of charities as you, it's easy to see how much good they are doing by your standards. Starting a charity or working for another is much more complicated because of the diversity of options. Depending on how pluralistic your values and epistemics are, inspiring others is more or less good.

This reasoning is analogous to another consideration, which is that doing builds more capacity and resources than researching. Research of this sort is relatively cheap to run, involving just the costs of salaries. However, running a direct charity requires far more employees and direct costs, such that one must build up a larger donor network to run it. For example, to run our research program costs \$25,000 USD this year, whereas to run Charity Science Health will cost \$250,000 for the first year, and could well reach the multiple million dollar per year mark. On the other hand, if the cause you end up choosing is very different from your initial top option, then many of the donor network you built up for that first charity may not be interested in your next choice.

Nonetheless, this ends up not being too large of a consideration, because building up resources likely follows an s-shaped curve. This means that even if you start a few years later than your counterfactual self, you will eventually more or less catch up with them in terms of resources.

Learning by doing

There's a great quote from **Brian Tomasik**: "There's a quote attributed to Abraham Lincoln (**perhaps incorrectly**): "Give me six hours to chop down a tree, and I will spend four hours sharpening the axe." This nicely illustrates the idea of front-loading learning, but I would modify the recommendation a little. First try taking a few whacks, to see how sharp the axe is. Get experience with chopping. Identify which parts of the process will be a bottleneck (axe sharpness, your stamina, etc.). Then do some axe sharpening or resting or whatever, come back, and try some more. Keep repeating the process, identifying along which variables you need most improvement, and then refine those. This **agile** approach avoids waiting to the last minute, only to discover that you've overlooked the most important limiting factor."

This covers a rather important advantage of doing – that you're not just doing. Learning never stops. How much should we take this into account?

I think that this is definitely important because it helps determine which options are realistic and helps calibrate your probabilities. Indeed, historically there have been many things that maybe I could have learned via research, but I probably wouldn't have without getting my hands dirty.

On the other hand, learning by doing is learning by anecdotes. Learning through reading is learning through thousands of anecdotes, otherwise known as science, or even just picking up the individual anecdotes of many others. Additionally, there are some things that you can simply never "learn by doing", which includes many crucial considerations. For example, you can't just work for a charity and naturally pick up which is better, frequentism or Bayes, or whether you should be speciesist or not. Those are things that need explicit reasoning and research.

Furthermore, learning by doing is very costly per amount learned compared to direct learning. Getting a job or starting a project in an arena is a huge investment which is hard to pull back from once you've started.

On the other hand, you can risk losing touch with reality if you do not have some hands-on experience. It also lessens the gap between learning via research compared to implementing your top option.

Fortunately for me I share an office with a direct implementation organization, so I get the benefits of both worlds, and I have not felt the need to completely address this question. This may be hard to replicate, but some alternatives, like befriending those doing direct work, might confer similar benefits.

Value of research

The value of research is, rather straightforwardly, the increased value of your best choice. A great example is that when I started my altruistic career as a child, I saved [kelp](#). My grandmother had told me that kelp were alive. I took this to mean that it was sentient, and then spent many hours in the summer saving kelp from drying out on the beach and having a painful, drawn out death. In retrospect this was adorable, but 0% effective. Given a vast increase in knowledge since then, I have since learned that kelp are not sentient, and given my increased understanding of the world, am helping people at a much larger scale. The value of my best option increased enormously.

The key question then is how much the marginal amount of research increases the value of your best option. This is impossible to answer precisely because we'd need to know the end result, and if we did, we wouldn't need to do the research. Fortunately we have a way to deal with uncertainty in this domain, which is the expected value of information. [Peter Hurford has a great post on this](#) which is generally the method I followed. I just added the concepts of remaining years left to figure out how it compared to doing.

Which brings me to the last concept: that you should switch from researching to acting once the expected value of marginal research equals the expected value of acting. The expected value of research will go up for a while, then start going down as you've thought of most of the relevant considerations. It will also start going down in life as you have less and less time to be able to capitalize on this knowledge, which will eventually nudge you into action.

To calculate the marginal value of one additional year spent researching, you can follow this formula:

$$[(\text{Change in value of best option}) \times (\text{Percentage of value added by added year researching}) \times (\text{Working years left} - \text{Years spent researching})] - (\text{Value achieved if you researched one less year}).$$

Simplified this is just the value of $t+1$ years of research minus the value of t years of research.

Calculate this for each year until the calculation gives a number less than 0, at which point switch to doing.

Of note, in this model, I assume that the percentage of value added per year of research is a consistent percentage of the remaining value. This means you get closer and closer to 100%, but never there. So if I expect a value increase of 10 times if I researched forever, and to achieve 50% of the value for each additional year of time, I would expect to get 5x the value the first year, then 50% of the remaining 5, so $50\% \times 5 = 2.5$ the next year, etc.

Here's a worked example:

Expected change in value of best option = 5 times better than current option

Proportion of remaining potential value for each marginal year of research = 70%

Working years total = 40

$[70\% \times 5 \text{ times better} \times (40 \text{ years} - 1 \text{ year researching})] - 40 \text{ years at current value if just research} = 96.5$. This is positive, so try the next year.

$[((5 - 3.5) \times 70\% + 3.5) \times (40 \text{ years} - 2 \text{ years researching})] - [70\% \times 5 \text{ times better} \times (40 \text{ years} - 1 \text{ year researching})] = 172.9 - 136.5 = 36.4$. This is positive, so try again.

$[((5 - 4.55) \times 70\% + 4.55) \times (40 \text{ years} - 3 \text{ years researching})] - [((5 - 3.5) \times 70\% + 3.5) \times (40 \text{ years} - 2 \text{ years researching})] = 180 - 172.9 = 7.1$. This is positive but close to 0, so we're getting close.

$[((5 - 4.865) \times 70\% + 4.865) \times (40 \text{ years} - 4 \text{ years researching})] - [((5 - 4.55) \times 70\% + 4.55) \times (40 \text{ years} - 3 \text{ years researching})] = 178.5 - 180 = -1.5$. This is negative, but just barely, so it indicates that you should spend a little under 4 years researching before moving on to acting.

Of course there are many limitations to this calculation. The three main ones are:

- Immense uncertainty. Each of the parameters are best guesses, but "guess" is a key term. How can you tell whether you're at 70% of the value of research or at 10%? Likewise, how many working years do you have left? These are all **highly uncertain so cannot be taken literally**, but they're **better than no estimate whatsoever**.
- Time consuming. This might be a consequence of my method, and I'm sure there are more elegant ways to calculate this.

- Simplifying assumptions. The working years left included delegability, but it didn't take into account the nuances of it. It might complicate the formula quite a lot to take into account that some delegation can be done while researching, whereas others you have to take time off of research to get rolling.

We ran these calculations with a variety of optimistic, pessimistic, and best-guess scenarios and all of the results came out in the 2 to 8 person-year range. The next question is what to do with these numbers. Two to eight years is a wide range and the numbers are uncertain thus subject to wide fluctuations based on new information. Our conclusion has been to follow the general process of:

- List all possibilities. Make a list of all crucial considerations and possible causes to investigate.
- Divide into chunks. Divide the average time from the calculations above into equal chunks for the crucial considerations and causes, based on previous such research. We came to about 1 month on each consideration/cause.
- Divide chunks in two. Divide those chunks in two, in this case 2 weeks. Do an initial run through of all the crucial considerations, taking 2 weeks on each one.
- Allocate remaining based on need. Based on the progress and new information from this initial passover, budget the aggregated 2 weeks per consideration left to the most promising/in need considerations.
- Redo calculations. At the end of this reassess and redo the calculations based on all of the new information. Spend longer or less time on crucial considerations based on this.
- Repeat. Do the same for cause comparison.

The advantages of this method compared to others considered are that it's time saving, with deadlines making you work faster, and the benefit of seeing things with fresh eyes. It makes sense too because the calculations are only rough approximations, so do not give enough precision to make day-to-day decisions in any case.

Spaced research throughout rest of life

This is half the puzzle. You cannot simply research once and then call it a day. The world changes and there will be new considerations. Thus part of the solution is to do spaced out research phases throughout the rest of your life. So, how should they be spaced out? We've decided to postpone on that decision until after the initial phase of research, but here are some contenders we thought of:

- Sabbatical model, where once every few years you take a year or a few months off to incorporate new considerations.
- Vacation model, where you take multiple one week long “vacations” per year to do research.
- Project based, where you start a charity/project, stay with it until you can step back without hurting it, then do another round of research between projects.
- Progressive investment in delegation, where you start with low energy levels of trying to inspire, say posting a single blog post on the topic, then progressively intensify your energy based on interest. If somebody starts a project based on the blog, then go back to research. If nobody does, try more actively advocating it. If that doesn’t work, keep investing more until you get to the point of starting it yourself.
- Needs basis, where you take time off on a needs basis.
- Completely passive, where you simply do the research in your spare time after the initial upfront investment.
- Split your time. Always devote a certain percentage of your working time to doing research.
- Spaced repetition, where you space out your research closer together at the beginning, then space it out further and further as time goes on and you’re more confident in your beliefs.
- Always have an “R&D” division of your organization, where there’s always at least one person who’s full-time thinking about what to do next, while there are also people executing on the current best plan.

Our current best guess is a combination of progressive investment based on delegability and project-based. This seems to take advantage of delegability without having to rely on it entirely.

Remaining questions

These considerations are currently incomplete. Some of the weaknesses we plan on investigating further as separate crucial considerations, some we might come back to if we think optimal stopping has been stopped at the optimal time or not. These gaps include:

- **Knightian uncertainty.** What should we do when there is enormous uncertainty around an issue? Is it better to give your best guess at a number, or is it better to follow a cluster approach? Maybe both or neither?

- Complexity of equation. The current equation is very simplified. It does not take into account different sorts of delegation. Is it worth making a more complicated formula to see how this affects things?
- Value percentage achieved per year. One key input into the formula is what percentage of the value of research you achieve per year. Currently we just add this based on a soft subjective sense. Is there any way we could improve our estimates?
- Learning by doing. The model we used does not take into account the benefits of learning by doing.

Ways to help

So there you have it. My current best thoughts on how long to spend researching versus doing. I would love your help. You can help by:

- Pointing out ways I could improve my reasoning, especially in ways that will lead to a changed conclusion. If it's just something that will reduce my confidence by a percentage point or two, that's less useful.
- Thinking of alternative approaches that beat my current one. This is generally a good way to help, because one's choice can only be so good as the best option you've thought of.
- Point me towards optimal stopping scenarios in the official literature that most apply to EA. I wasn't able to find any that corresponded that well with the EA situation, but that might have been a failure of thinking of the right keyword on my part.

3. How CE Uses Informed Consideration

Author:

Joey Savoie,

Co-founder & Director of Strategy

Charity Entrepreneurship's research process involves **multiple stages** and four methodologies: **cost-effectiveness analysis** (CEA), **expert views** (EpV), **weighted factor models** (WFM), and **informed consideration** (IC). In this article, we explain why and how Charity Entrepreneurship uses IC as part of our research process, and discuss its strengths and weaknesses as a methodology.

IC consists of using a non-predetermined and nondirected research process to evaluate an idea. It allows us to take into account forms of evidence and considerations that may be neglected by our other methodologies. It is particularly useful because some lines of reasoning are more theoretical or subjective in nature but still make up **crucial considerations**, and could thus make a charity idea worth pursuing or not. However, one drawback is that IC's output is generally formed by a single individual, which can significantly bias the results.

CE uses IC during three stages of our research at different levels of depth: first five minutes, then two hours, and finally twenty hours.

- First, we conduct an **idea sort** to sift out the top thirty or so ideas from our initial list of hundreds. We assess each idea using all four methodologies (i.e., CEA, EpV, WFM, and IC), spending twenty minutes per idea.
- Next we produce a **prioritization report**, spending two hours researching each idea from the idea sort and ordering them according to their promise. We use a different methodology for each cause area: IC is used for our **health policy** ideas. At this stage we only apply one methodology to ensure sufficient depth.
- Finally, our eighty-hour assessments of the top interventions involve all four methodologies.

Concretely, lead researchers use IC by doing broad research, trying to answer crucial considerations or intuitive questions related to factors that might dramatically

affect an idea's potential. They then assign an intuitive estimate of how promising a charity idea is.

What is informed consideration?


IC is the closest methodology we use to common sense, a priori reasoning, or intuition-based methodologies. It involves using a non-predetermined research process and takes into account soft forms of evidence and information that do not fit into our other methodologies. IC is generated by the lead researcher on a given report, and questions are asked that are more philosophical and broad in scope. It also allows space for undirected research that is not systematic in the same way that a WFM would suggest. It is described in a narrative form to help the reader get a sense of the information and intuition that the research team has on the issue after conducting substantial research.

Informed consideration example (Charity Entrepreneurship – 2019)

Size of the Animal

The second factor that plays a huge role is meat generated per animal or the size of the animal. Simply, cows are much, much larger than chickens. If you eat meat for a year you affect far more chickens than cows (even if you eat a perfectly equal weight in both). Chickens generate about 5 pounds of meat per animal, where cows generate around 750 pounds of meat. This results in the average person eating a lot more chickens than cows per year.

Within 10 years, the average American person consumes the meat of:



The results are that considerably more chickens are eaten by each person and consumed in the world as a whole. The below chart shows the very strong impact that chicken has relative to other land based animals. Broadly, the smaller the animal, the more eaten for the same number of calories, thus many more fish and chickens are eaten over cows or pigs.

Why is this a helpful methodology?

Many weak arguments can often be stronger than a single strong argument.

Sometimes an intuition can differ substantially from a calculated estimate even when both are based on the same information at the outset. Some lines of reasoning are more theoretical or subjective in nature but still make for crucial considerations that make a charity idea worth pursuing or not.

Reasons this is a helpful tool (in order of strength):

- Captures information that does not fit into other systems
- Captures information that would otherwise be lost
- Allows for divergent conclusions
- Accounts for many perspectives

- **Accounts for multiple reference classes**
- **Quick**
- **Flexible across ideas**
- **Formalizes existing intuitions and heuristics**
- **Generates cross-domain knowledge**

Captures information that does not fit into other systems: The biggest advantage of IC is that they take into account large pieces of information that are not easily captured by other systems. For example, maybe we know that a nonexpert group has a very strong view on an intervention, or there are some wisdom of the crowds or Chesterton's fence arguments in favor of an intervention. As formal models, including ours, do not offer a clear way to include this sort of information even to a small degree, IC gives both time and space to outline this sort of evidence.

Captures information that would otherwise be lost: Some information might fit into a different system but might not have been clearly captured. For example, if you have several quick conversations at a conference, you might get a sense of a given issue but not know exactly which individual or combination of individuals this sense came from. Information like this can easily get lost when creating deep formalized reports, but it is still important data and can lead to differences in researcher-intuition- vs. model-based results.

Allows space for more divergent conclusions: Many of the most valuable research contributions have come from divergent or original thinking. Detailed models do not often encourage divergent conclusions or thinking when applied to an intervention or issue. For example, IC would more easily spot if an intervention makes sense only when paired with a different intervention instead of being performed in isolation.

Takes into account many perspectives: Because IC includes both information that does not fit into other systems and information that would otherwise be lost, it ends up taking a large number of data points into account.

Takes into account multiple reference classes: One of the largest challenges in research is picking an appropriate reference class to compare the current intervention to. Often when comparing something it's easy to get stuck considering a single or limited number of reference classes, whereas softer comparisons often use a broader range of classes.

Is very quick: IC is a much faster system to use to go deeper on a specific issue. When there is not a prespecified question list or methodology, a researcher with the same amount of time can more quickly read and synthesize information than if they are seeking specific answers to specific questions. It is also one of the quickest ways to get a broad overview of a topic.

Has a high degree of flexibility across ideas: Some ideas require more time to be put into certain aspects compared to others. For example, an area that affects the total human population might require more philosophical consideration than a simpler intervention that merely makes already existing people less sick. Thanks to IC's lower level of structure, it can be used to direct more time toward idea-specific concerns.

Formalizes already existing intuitions/heuristics: One goal of written IC is to make more transparent the intuitions and heuristics that are used. Almost every important evaluative decision relies on intuitions or best guesses at key points, but having these considerations discussed and elaborated on makes it clearer where they come from, compared to a system such as a CEA where the full explanation might be "best guess."

Cross-domain knowledge is generated: Many times, hours used during IC are used for cross-applicable questions that might not be the most important questions for any single charity idea but are important for a wide range of charity ideas.

Why IC is not our only endline perspective

IC has a lot of breadth-based advantages; however, it is far from our only source of information and compared to the four other perspectives holds relatively minor weight in our endline perspective. Many of the same **concerns that exist for experts** also apply to IC but on a more dramatic scale. IC is generally formed by a single individual, making standard human biases a substantial concern.

- **Nonsystematicity**
- **Lack of precommitted methodology**
- **Nonnumerical**
- **Dependence on researcher strength**
- **Cognitive bias**
- **Overweighting personal experience**
- **Intractable issues**
- **Subjective individual judgment call**

- **Counting the same information multiple times**

Nonsystematicity: One of the biggest concerns with IC is that it is nonsystematic. Given the huge amount of data, conflicting experts, and possible angles from which to consider a complex issue, it is often easy for a researcher to find a long list of citations broadly supporting his or her views even if the views are wrong. IC is particularly susceptible to this concern.

Lack of precommitted methodology: The best studies conducted in science generally have a pre-analysis plan. There are **many resources** on why this is a good policy and improves research quality. Some forms of research have clearer ways to precommit to certain methodologies or sources. IC offers much more limited pathways than this when compared to more structured sets of questions for research used in the WFM.

Nonnumerical: IC is generally soft and nonnumerical, which can easily lead to factors holding inappropriate amounts of weight based on salience and general human weaknesses with numbers. It also makes many of the assumptions less clear because terms are softer such as “it seems probable,” which does not give the reader a clear sense of the percentage-based update or confidence of the claim.

Cognitive bias: IC has considerable likelihood of cognitive biases to the point of them being almost guaranteed. Some of these mental heuristics reflect useful but hard-to-trace-back knowledge; however, it is well documented that biases can consistently negatively affect decision-making in a wide range of ways, and IC is the most vulnerable system we use in this respect.

Dependence on researcher strength: IC leaves the bulk of the time distribution and research decisions to the individual researcher. This could lead to considerable quality differences between reports and researchers, making comparative results less useful. For example, a generally skeptical researcher might focus on flaws with these hours, and a more optimistic one might focus on implementation strategies. This could lead a reader or researcher to have an unreflective conclusion of what would have happened if a different research lead had been running the project.

Overweighting personal experience: Overweighting personal experience is a cognitive bias that is likely more severe for IC than for EpV. If someone has tried an idea or methodology and it did not work for them in the past, this is extremely likely to affect their judgment of similar ideas even when it should not. For example, imagine your odds of success at an activity are ten percent. If when you try you do

not succeed, your update should be almost nothing rationally but generally will be substantial intuition-wise.

Intractable issues: Given this methodology's focus on theoretical issues, it would be easy for a considerable amount of time to be spent on an issue that is not tractable. Some issues will be important but will eventually fall outside of the scope of an eighty-hour review. In addition, many issues that will be considered in this section cannot be moved forward even with several years of work.

Subjective individual judgment call: A final weakness of this methodology is that it leaves a lot more up to subjective judgment calls. This means that if your values or epistemology are largely different from those of our lead researcher, the results will be less cross applicable than the expert interviews or CEAs would be. It is also not easy to, for example, change a single input and see the difference in results in the same way as with the CEA or weighted quantitative models.

How much weight we give to informed consideration

We see IC as an important perspective but are also keenly aware of its many documented weaknesses. It is one of our four methods, but we expect it to hold between one-fifth and one-tenth weighting, with considerable variation depending on the specific charity idea and cause area. This is far less than CEAs, EpV, or WFM, beating only prior views (our fifth perspective). We expected IC to be comparatively stronger in areas where there are many philosophical considerations with an unclear consensus on the answers.

How CE generates informed consideration

IC has quite a few unique elements compared to our other systems.

Broad research

In theory: We have written about the use of **broad research** when understanding a cause area in the past. Broad research in this context is a similar concept but applied to a more specific intervention. For a first sense of an intervention, it is easy to rely on a few sources or types of sources for the bulk of your information. For example, it would be easy to read many studies and books written by academics but gather limited information from field workers who often have a very different perspective. The point of broad research is to, somewhat slowly, get a sense of the area from multiple angles.

In practice: This involves watching related documentaries, reading related books, going through overview sources such as Wikipedia, and trying to look at the issue from as many different introductory perspectives as possible. The exact way that hours are used in broad research will vary depending on the cause area and where the most information can be found but often involves many different methods and sources.

Example: When conducting broad research on malaria-preventing bed nets, one might read a book on the historical progress of the fight against malaria; watch a short documentary on a specific village that is combating malaria; read the Wikipedia pages on malaria, bed nets, and other related interventions; read blog posts from half a dozen different authors writing on the topic, and listen to a podcast interviewing a bed net producer.

The main questions to consider in the broad research phase include:

- Have I looked at all the related summary material that is available on evaluation websites, CE's website, or cross-cutting organizations?
- Have I considered at least four different clusters with different views on this issue and read some resources from each?
- Have I reviewed information from a number of different media (e.g., studies, news articles, Wikipedia, documentaries, interviews)?

Undirected research

In theory: When IC is compared to our other methodologies, it is clear that it is the least directed. The research questions are far less detailed, and much is left up to the researcher's interpretation of what to prioritize and how many hours to spend in each area. This reduces the systematicity of the intervention but opens up the ability to spend more hours in key areas that were unresolved from the more systematic research conducted.

In practice: This means there are a few guiding questions for the IC write-up and an expected consistent end product. However, the specific sources and hour spread will vary considerably between cause and interventions, and cross-area uniformity is far less expected.

Example: During the broad research phase, a key consideration is found; it seems this single consideration could determine whether the intervention is impactful or not. The researcher decides to spend the bulk of his or her time deep-diving this

consideration because it seems both trackable and pivotal to an eventual charity recommendation in the area.

Questions to consider:

- What are the factors least likely to be considered deeply enough by our other methodologies?
- What types of research have been leading to the best progress on this charity idea?
- What issues do I feel like our team needs to consider more?

Crucial considerations

In theory: Crucial considerations are considerations that could significantly affect an intervention. In some cases a consideration will be both important and cross-cutting throughout many charity ideas and thus will be worth an independent supporting report. However, in many cases an issue will be worth more consideration than our systems give it but less than a full independent report.

In practice: When researching an intervention, a consideration that is very important to the conclusion will present itself that will require dedicated hours. Often it will be worth trying to brainstorm multiple possible crucial considerations and determining an hour split between them (depending on their importance and expected progress on them) before diving into one.

Example: When considering an online cognitive-behavioral therapy (CBT) application, one of the biggest questions might be whether online therapy has evidence of working and, if so, in what contexts. This issue is likely large enough to be worth spending some unstructured research hours on but not large enough to warrant a separate supporting report.

Questions to consider:

- What are the top crucial considerations in this area?
- How likely am I to make progress on this consideration given the hours I have?
- Does this affect enough charity ideas to be worth a separate supporting report?

Theoretical considerations

In theory: Certain questions require more introspective or philosophical consideration to make progress on. Writing down your thoughts can be a way to

mediate this sort of thought in a more considered and transparent way. Chains of logic **can sometimes** lead to surprising conclusions when given room for consideration.

In practice: In practice, some questions will be ethical or philosophical in nature, and there will not be a clear answer. Many relevant questions will not have written resources to rely on, and the researcher will have to determine how to make progress on the issue.

Example: Is a smaller amount of harm spread across a larger number of individuals a permissible trade for a reduction of a large amount of harm spread across a smaller number? At what point do most people agree with the trade-offs, and at what point is there controversy? Could I conduct a quick survey to get a sense of these numbers?

Questions to consider:

- What ethical views would change our conclusion on this?
- Are there any jumps in logic or reasoning in our research?
- What evidence could falsify or change my conclusion on this?

Intuitive considerations

In theory: Often intuition can pick up on things that more formal systems might miss. It can be affected by anything from heuristics from other areas to subtle body language that an expert had but did not make it into the notes.

In practice: We have intuitive feelings about everything that constantly affect our judgment; making these feelings and views explicit and trying to see if they are relevant information or not is an important part of the process. We can also flag for external reviewers what areas might be affected by intuition.

Example: A certain number in the CEA result just seems far too high; it seems to affect and trivialize everything else. I should really research that key number more to make sure it holds up on a double check.

Questions to consider:

- Does my intuition feel uncomfortable with any of the conclusions reached on this report? If so, why should I put more time into that spot?
- Is there any area or question I am intuitively shying away from?

- Could I take some of these conclusions further than I currently do?

Different lengths of informed consideration-based estimates

Five minutes

Five minutes is the length of time we spend on each methodology during our initial **idea sort**. It is not enough time to do even the smallest amount of broad research for each charity idea; however, it is enough time to do a minor amount of broad research on the cause area as a whole and then make an intuitive estimate on how promising a charity idea is. For example, if there are 300 ideas in a space, twenty hours could be spent doing unstructured broad reading in the cause area. You do not need to take polished notes on this research because the only published product will be the list of ideas. Afterward, work through the list spending one minute on each idea. For each idea ask yourself a question—“What is the probability that we will recommend this as a charity idea?”—and rate it based on that.

Broad research questions to consider:

- Has CE written any reports in related areas?
- Have any charity evaluators (GiveWell, Animal Charity Evaluators, Copenhagen Consensus Center), cross-cutting organizations (Rethink Priorities, Wikipedia, World Health Organization, Abdul Latif Jameel Poverty Action Lab (J-PAL), Institute for Health Metrics and Evaluation), systematic reviews and meta-analysis organizations (Cochrane), or think tanks (Center For Global Development, The International Initiative for Impact Evaluation (3ie), Department for International Development) written any overviews of the topic?
- Who is considered the respected source in the area? What are they focusing on, and what have they written on the topic?
- It is up to the researcher’s discretion what questions are going to bring most value to inform their judgment.

Expected outcomes

- A score out of ten for each of the ideas

Advice:

- Before rating the idea, **do a calibration exercise**
- Organize your twenty hours beforehand to make sure you spend enough time on important considerations. Spend time listing your biggest questions and try to make sure you keep enough time for each. An example of such an exercise can be found below. Process:

1. Set up the time cap
2. Compile list of possible sources
3. Guess how useful they might be on a scale of zero to ten
4. Estimate how long it might take to read each source
5. Order by expected usefulness and sum the expected time to decide which sources to read (best ratio of time to usefulness)
6. Read all of those sources within the time cap

	A	B	C	D	E
1	Source	Expected usefulness	Estimated time (minutes)	Time tracked (minutes)	Link
9	Trends in Contraceptive Use Worldwide	10	20	15	https://cl.ly/ff508d83d47c 1) At least one in ten married or in-union women in most regions of the world has an unmet need for family planning. 2) Contraceptive prevalence is projected to increase from 17 to 27 per cent in Western Africa, from 23 to 34 per cent in Middle Africa, from 40 to 45 per cent in Eastern Africa, from 55 to 60 per cent in South Asia, from 60 to 65 per cent in Southeast Asia, from 65 to 70 per cent in Latin America and the Caribbean, and from 70 to 75 per cent in Europe and Central Asia. 3) Female sterilization and the IUD are the two most common methods used by married or in-union women worldwide used rhythm or withdrawal. There are large regional differences in the use of some types of contraceptive methods. 4) https://cl.ly/977fe519e040 5) https://cl.ly/2f0a650a10f8 6) https://cl.ly/c9f0c8b3802d
10	FAOSTAT: Food balance - food balance sheets	10	5	5	https://cl.ly/7fd7f2cc7584
11	2017 The state of food and agriculture: Leveraging Food Systems for Inclusive Rural Transformations	7	20	15	https://cl.ly/5f8b7ced3ee0 1) Shares of income from different sources in rural areas of developing countries https://cl.ly/7fd7f2cc7584 2) Rate of annual population growth in larger cities and in small cities in selected countries https://cl.ly/5f8b7ced3ee0 3) Distribution of global population along the rural-urban spectrum https://cl.ly/5f8b7ced3ee0 4) Distribution of global population along the rural-urban spectrum globally and by region https://cl.ly/835dc3d1 5) Ratio of rural to urban population for smaller cities, larger cities and rural, globally and by region https://cl.ly/4 6) Meat, fish and dairy consumption increases strongly with income growth. 7) In general, the shift happens earlier in urban than in rural areas and, as Bennett's law predicts, increases as income rises. 8) Complementing the trend toward a reduction in the share of cereals, there is a shift to non-staple foods – especially fruits, vegetables, meat, fish and dairy. 9) The daily per-capita consumption of protein from animal sources in low- and middle-income countries rose from 1960 to 2010. Data on Malawi, Uganda, the United Republic of Tanzania, and Zambia, especially for the consumption of meat, fish and dairy, show a clear upward trend. This suggests that the easier accessibility of these products is a key driver of the trend. 10) The study by Reardon et al. (2014), covering Bangladesh, Indonesia, Nepal and Viet Nam, showed that as income rises, the share of animal products in the diet increases. Meat and fish accounted for an average of 30 percent of the urban budget in South Asia and Southeast Asia in 2007. In three decades, the domestic aquaculture supply chain in Bangladesh (see Box 4) grew in volume 25 times. The outcome for the food system of the above dietary changes is the very rapid and massive rise of supply of animal products. 11) Fisheries and aquaculture are an important source of food, nutrition, income and livelihoods for hundreds of millions of people. Incomes will see strong growth in markets for fish and fish products in the next decade (HLPE, 2014). In China, the share of animal products in the diet is rising rapidly.
12	2009 The state of food and agriculture: Livestock in the balance	8	20	23	https://cl.ly/c47fad2cbb6 1) Rapid growth and technological innovation have led to profound structural changes in the livestock sector. It has increased the availability of livestock products and has increased rapidly in developing countries over the past few decades. 2) Trends in consumption: Consumption of livestock products has increased rapidly in developing countries over the past few decades. 3) https://cl.ly/c47fad2cbb6 4) Consumption has increased in all regions except sub-Saharan Africa. Also, the former centrally planned economies have seen a decline in consumption. In South Asia, including India, per capita consumption of livestock products has increased significantly. 5) Consumption of livestock products per capita in developing regions is still substantially lower than in the developed world. 6) Drivers of consumption growth: The growing demand for livestock products in a number of developing countries has been driven by economic growth, rising per capita incomes and urbanization. 7) Developing country meat exports are dominated by the contribution of Brazil, the world's largest meat exporter. 8) All indications are for continued growth in global demand for livestock products. In 2007, the "IMPACT" model projected that global demand for livestock products would increase by 100 percent by 2050.
13	Livestock commodities: Past and present	5	5	2	https://cl.ly/7fd7f2cc7584 1) Milk and dairy products, production and use: past and projected 2) Table 3.10: Food consumption of meat 3) Table 3.12: World exports of livestock products and percentage of world consumption 4) Table 3.13: Net trade positions of the major importers and exporters of livestock products
14	SUM MINUTES		160	138	
15	SUM HOURS		2.666666667	2.3	
16					
17	Meat Consumption Statistics	7	5		https://cl.ly/7fd7f2cc7584
18	2012 The state of food and agriculture: Investing in agriculture for a better future	4	15		https://cl.ly/7fd7f2cc7584
19	FAO Global livestock production and health	4	10		https://cl.ly/7fd7f2cc7584
20	Change in the livestock sector	4	10		https://cl.ly/7fd7f2cc7584
21	OECD-FAO Agricultural Outlook: every year different focus	4	20		https://cl.ly/7fd7f2cc7584
22	PMA 2020 annual surveys for family planning in eleven countries	8	20		https://cl.ly/7fd7f2cc7584
23	2007 The state of food and agriculture: Paying farmers for environmental services	3	15		
24	FAO CountrySTAT	4	10		https://cl.ly/7fd7f2cc7584
25	FAO Aquaculture and Fishery Stats	5	10		https://cl.ly/7fd7f2cc7584
26	FAO 2005 The state of food and agriculture: Agricultural trade and poverty	5	30		https://cl.ly/7fd7f2cc7584

Two hours

In the two-hour model, the researcher has enough time to do unstructured research on the given charity idea. Some of this time will be spent finding sources – ideally overview sources – and some of the time on reading and absorbing them. This knowledge will build upon the twenty hours of cause area knowledge that has already been acquired. The bulk of this time will be spent on unstructured research with the last roughly thirty minutes spent on contemplation and written description

(one paragraph) of why an idea is ranked more or less promising. This stage will only happen for one of the four cause areas we are researching.

Questions to consider:

- What factor might rule out this intervention quickly?
- Are there any overview books/deeper write-ups that might summarize the area (e.g., a section of a textbook)?
- Are there specific write-ups or blogs that focus on this area? Can you read a sample of them to get a sense of the broader space?

Expected outcomes

- One paragraph on why the idea was ranked more or less promising

Twenty hours

IC happens at the start and the end of the eighty-hour report research process, with roughly ten hours spent at each point.

- At the start, IC is used to attain a deeper view of both the broad research and the key **crucial considerations** that might dramatically affect the intervention. This will take about ten hours of the IC process and will result in a one-to-five-page summary of the critical considerations, mitigations, and overall feel of the intervention.
- Near the end of the research report, the researcher contemplates all the data they have obtained over the course of the research process. Part of this **internal contemplation** involves writing summaries and updating learnings from all the other areas as well as giving some time to introspective thinking and theoretical considerations that need further work. The expected output is also one to five pages of writing. This section can be less structured and less cited than other parts of the full report and should include intuition and gut checks on various assumptions, even if it is unclear where the intuition comes from. Internal contemplation is also the chance to include points that might not be clear from the other models but would be clear if someone conducted the full scope of research themselves.

Questions to consider

- What would every person in this field have read or heard about?
- Consider if there are any crucial considerations or other factors that could lead to elimination of the idea. **Steelman** back and forth to see if the elimination criteria are strong enough.
- What are the predominant considerations against doing this intervention?

- What are the predominant considerations for doing this intervention?
- How does it seem from the reading you have done, past experience, and people you have talked to?
- What do we not know?
 - What further research should be done if we were to go deeper with this approach?
 - What are the open or remaining questions on the issue?
 - How much more research might it take (if any) to feel confident?
- What do experts seem to broadly think of this area?
- How does this intervention broadly compare to the other intervention areas?
- Conduct pre-mortems:
 - Imagine it's a year from now, and this intervention has been tried and has failed.
 - Write down why you think it failed, considering both internal and external factors.
 - What actions could you take to prevent this hypothetical failure mode?
 - Imagine that you apply them.
 - Iterate on the above steps, imagining that even with your new failsafes, the plan still fails. Are you shocked? How much?
 - Do the same for a time horizon of five years.

How IC compares to other methods used (timeline of an 80-hour report)

- Ten hours – Broad undirected reading and CC (IC)
- Sixteen hours – Directed research (WFM)
- Ten hours – Finding and talking to experts (EpV)
- Twenty hours – CEA creation (CEA)
- Four hours – Directed research (WFM)
- Ten hours – **Summary writing** and internal contemplation (IC)
- Ten hours – Showing endline report to experts (EpV)

Expected outcomes

- First section: two-to-five-page summary of crucial considerations considered and resulting research
- Second section: one to three pages of internal contemplation as well as the writing of the summary report

Summary

By the time a charity is recommended, it will have been thoughtfully considered multiple times at different points in the process. Crucial conditions and theoretical

questions will have been researched; broad and indirect research will have been carried out. In many ways this part of the process is most similar to how the bulk of research is conducted, with few formal systems set up beforehand and the lead researcher's endline views carrying a lot of the conclusion's weight.

Reading list

Internal resources

- The benefit of broad understanding
- Iterative depth
- Other methodologies used in our research:
 - Cost-effectiveness analysis
 - Expert views
 - Weighted factor models

External resources

- When can you trust your gut?
- Sequence thinking vs cluster thinking
- Importance of crucial considerations
- Many weak arguments
- Sources of Power: How People Make Decisions
- The Power of Intuition, and Streetlights and Shadows: Searching for the Keys to Adaptive Decision Making
- Thinking, Fast and Slow by Daniel Kahneman

IV. Tools of science (the screwdriver)

Science loves taking things apart and understanding how they work. Often scientific progress is slow and incremental, like slowly tightening a screw. Matching the hundreds of types of screwdrivers, there is a huge range of scientific techniques that are useful in a wide range of different situations. And sadly in both cases, there is limited standardization in the methodological process.

Science is likely the most established and proven of the tools that we cover here. While you will not be able to learn it at the same speed as rationality, it can be used to build a much stronger foundation of evidence and confidence in the way the world works. In this section, we will cover:

- How convergence and robustness of data can lead to confidence in its conclusion
- Falsifiability and how to make sure your charity can be tested
- Different types of scientific evidence and the evidence pyramid
- Tricks people play that can make scientific evidence misleading

1. The Importance of Evidence

Author:

Joey Savoie,
Co-founder & Director of Strategy

Any money spent on an intervention that doesn't work is money wasted. Money that could have saved lives. This is why evidence is so important – you've got to have a high level of confidence that the intervention is having the intended effect.

Imagine two charities: Stop AIDS (a made-up charity for illustration's sake), and Homeopaths Without Borders (a real charity, unfortunately). Stop AIDS provides antiretroviral drugs, which are clinically-proven to dramatically improve the quality and quantity of life of HIV positive patients. Homeopaths Without Borders provides 'medicine' that has been proven **ineffective**.

Now, we chose an obviously ineffective charity to prove our point, but there are many other charities that have no scientific evidence of their impact. In fact, it's **the norm**. And there are a fair few charities out there that have been evaluated and deemed harmful, yet they're still actively seeking and accepting funding.

So, how can you identify good evidence of effectiveness? Your best bet is to look for controlled studies, with an emphasis on falsifiability and a convergence of evidence.

Look for convergence

Convergence is a very important but oft-forgotten concept. In short, the more diverse pieces of evidence that converge on the same conclusion, the more confident we can be about that conclusion. This is because if they are indeed independent sources of evidence, it's unlikely that they would converge on the same conclusion unless it were true.

Convergence suggests robustness – the idea that a claim holds true even when you look at it from a different angle. Such ambiguities ought to concern any charity founder. How can we be sure some intervention will work in one place, just because it worked somewhere else? Many people place too much weight on one strong study, changing their minds too rapidly, even if previous weaker studies or macro-level data point the other way.

Falsifiability

What is falsifiability?

A claim is 'falsifiable' if there is evidence that could prove it does not work. In other words, we can make a prediction and test it, and if the prediction is inaccurate, then the claim is proven to be false. However, if a claim is 'unfalsifiable', that means it can never be proven or disproven.

How can you spot an 'unfalsifiable' intervention or charity?

'Unfalsifiable' charities ignore any evidence against their effectiveness or explain it away with post hoc rationalizations.

For example, many studies have shown that **microloans don't increase income**, and it's **better to give the poor money** with no strings attached. However, microloan charities still exist and either ignore this evidence, or say that the real point of

microloans is to empower women. When more evidence comes out casting **doubt on whether they are even successful at that**, they say the point of their services is actually to smooth out income fluctuations. This claim might be true, but at this stage you should smell a rat – you’d need to wait to find evidence supporting this, and it’s starting to look like microloans are not all they’re cracked up to be.

Note that microloans are technically falsifiable, but psychologically not so. They did indeed make claims that were disproven. Nonetheless, the proponents either refused to see the evidence, or moved the goalposts, making it essentially unfalsifiable.

What’s the problem with unfalsifiable charities?

Just as scientists had to accept that the world is round, charities must learn to admit when they realize their interventions aren’t working. We can never improve the world if we refuse to acknowledge and learn from our mistakes.

How can you make sure your intervention is falsifiable? There are four main ways to hold yourself accountable to the evidence:

1. **Make sure you can clearly define success and failure**, so you cannot weasel out of failures later by saying you were trying to accomplish something different anyway.
2. **Counter confirmation bias** by actively trying to prove your intervention wrong.
3. **Be prepared to abandon or change your intervention** if it doesn’t work. It helps to have a **Plan B** you’re excited about, so that you do not feel like the world will end if Plan A doesn’t work out.
4. **Get an external review**. It's very hard to evaluate the evidence objectively, but an intelligent third party can be much more impartial than you. (GiveWell is a good example of this.)

An exemplary falsifiable charity: GiveDirectly

Take **GiveDirectly** as an example of a charity that's as falsifiable as they come. They give money to the poorest people in Kenya and Uganda, no strings attached. There are **11 randomized controlled trials** proving that this intervention improves long

term income, empowers women, and makes people happier, which we believe is the ultimate metric for a charity's success.

Not only that, but one of the studies was run by GiveDirectly themselves, and they publicly defined what they would regard as success or failure, explained how they would analyze the data, and precommitted to publishing the results regardless. This was a bold and admirable move. They could easily have found that the data did not confirm their intervention's success.

You can see their dedication to transparency even in their marketing. In one of their social media campaigns, they told stories about what people spent the money on. Unlike the vast majority of other charities though, they did not cherry-pick the most compelling stories. They randomly selected them, so the examples were truly representative.

GiveDirectly is one of the very few falsifiable charities that goes to great lengths to test itself. That means you can trust them to actually do what they say they will.

Commonly used types of evidence

There are thousands of different types of evidence ranging from casual anecdotes to randomized controlled trials with hundreds of thousands of participants. How can you know which sources to trust when researching interventions? We summarise some of the key pros and cons of the types of evidence charities use most commonly, in order from strongest to weakest.

Meta-analyses and systematic reviews

Meta-analyses and systematic overviews are at the top of the evidence pyramid. These involve systematically identifying all the studies that satisfy predefined criteria, such as strength of study or target population. They then analyse the dataset and determine which conclusion is best supported by all the currently available experiments. General standards for meta-analyses and systematic reviews make them less prone to bias than simply perusing the existing literature according to interest and coming to a conclusion at an arbitrary point.

This is a huge part of the reason we like the work done by organizations like GiveWell, as they can look at large numbers of studies and put together important overviews that allow others to easily and painlessly come to an informed opinion.

Individual scientific studies

Generally, scientific studies are very reliable sources of information. In the medical sciences, studies are necessary to determine whether a new drug is working. For charities, studies are equally important to prove whether the charity is having a real positive impact. However, studies can vary widely in strength from randomized controlled trials (RCTs) to much weaker observational studies. Studies vary in quality, so do not take the results at face value, but rather evaluate each study based on its own merits.

Expert opinions

The reliability of ‘expert opinions’ is entirely dependent on who the expert is and the field they are in. For example, a recommendation from a prominent expert in a relevant field is usually much more trustworthy than a recommendation from an academic who specializes in another subject.

You must also try to gauge whether the expert is biased. For example, academics tend to favor their own fields and their own work within them. A researcher who specializes in vitamin A supplements may think that their cause area is the most important and another researcher in education will think the same, though both do not have much understanding of the other’s field. While they may know within their area which is the best intervention, they will be hard pressed to compare between areas.

This natural bias can be used to your advantage, however. When considering whether to implement an intervention, try running your idea past several reasonable outsiders whose values and epistemologies you respect. If many of them give you negative feedback, you can be pretty sure that the idea isn’t as good as you thought, because the vast majority of people tend toward over-optimism, or at least try to avoid hurting your feelings by being positive.

Additionally, some subjects give better expert advice than others according to how tractable the subject is. A mechanic, for instance, will likely know exactly what to do to fix a car. Cars have limited moving parts and each part is understood well, with clear and quick feedback loops. On the other hand, if you ask ten futurists what will happen in 50 years you will get ten different answers. There are countless moving parts, each of which is poorly understood, and feedback loops are muddy and slow. While most of the charitable areas you will likely consider do not have the advantage of having such a well understood field as car mechanics, some fields still do have more reliable experts than others. For example, health experts are more

reliable than experts on social change, which is a much messier and less well understood field.

Be aware that, while your source may have some valuable insights, it's also equally possible that his or her recommendation could be worthless. Even if you think your recommender is genuinely reputable, you should be sure to question their original sources. Is their opinion based on personal stories and reasoning, or more solid evidence like studies and systematic overviews? If at all possible, rely on the primary sources that the expert is using rather than the opinion of the expert themselves. Expert opinion should generally be used as a jumping off point or as a resource if you are not planning on putting much time into the area, rather than a primary source for choosing a charity to start.

Historical evidence

Some people use historical evidence to predict which interventions will work and which might require more thorough testing. Sadly, there are significant problems with most historical evidence. In practice, most charities' historical evidence is largely unreliable and cherry-picked to support a predetermined conclusion.

Some general characteristics of good and bad historical evidence are listed below:

Bad historical evidence	Good historical evidence
Theory before evidence	Evidence before theory
Non-systematic (researcher follows interest)	Systematic (researcher follows a predetermined process)
Anecdotal and cherry-picked	Large, representative sample
Strong causal attribution	Cautious causal attribution
Only looks at successes	Looks at successes and failures
Shallow research (examined few resources)	Deep research (examined many sources)
A bad example: "Protesting worked for the civil rights movement in the past, so it should work for our cause!"	A good example: <i>GiveWell's history of philanthropy</i>

A priori reasoning

A priori reasoning is based on logic or common sense. For example, “Women in developing countries have to collect water at pumps, which is long and difficult. Children have a lot of pent up energy but no playgrounds. If we make a merry-go-round that pumps water, the children will have a place to play and run around while simultaneously getting water and the women will be free to follow other pursuits.”

Logical or common sense reasoning can feel more reliable than personal stories, but it suffers from many of the same problems. The reliability of a priori reasoning depends partly on the field of research. For example, reasoning in mathematics is more useful than historical evidence, but historical evidence in sociology is more useful than a priori reasoning.

In the charity sector, the world is messy and rarely goes according to plan. For that reason, you should not rely on this kind of evidence alone to evaluate charities and interventions. When lives are at stake, armchair reasoning is not enough.

Anecdotes and personal stories

Anecdotes and personal stories are the weakest forms of evidence and, sadly, the most commonly used in the charity sector. They are almost always cherry-picked and unrepresentative of the norm. You should never rely on them as your sole form of evidence for an intervention’s effectiveness.

However, that’s not to say that personal stories don’t have a place in the charity sector. Good personal stories can be extremely compelling emotionally, and they can be a great way to attract donations for a cause that other evidence already supports.

Getting the evidence yourself

It’s often a good idea to run an RCT on your own intervention, even if there’s already strong evidence that it has been effective in the past. You might be running a slightly different intervention or working in a different context than previous experiments. You may also want to measure metrics that other RCTs do not typically measure, such as subjective well-being, or look into any flow-through effects that may concern you.

The case for testing new ideas without any evidence – yet

Under the right circumstances, testing new and innovative ideas can be very valuable. For example, testing allows you to share your results with other charities, which can then learn from your mistakes or replicate your successes. Without testing new unproven concepts, we may never find the most effective interventions. And in this circumstance, settling for less than the ‘best’ intervention means accepting the fact your target population will suffer more, or for a longer period of time.

Of course, not every idea is worth trying and we wouldn’t usually recommend taking a chance on an intervention with zero evidence to suggest it may work. We believe the most promising interventions for trial and error have some prior evidence suggesting they may be important, but not enough evidence to answer all important questions. If the intervention has a foundation of suggestive positive evidence then the chance that your research will turn out to be totally useless or irrelevant is lower.

Time

Weigh the existing evidence for the innovation against the resources required to test it. For example, if an idea only takes a day to test, we require much weaker evidence to try it than if the test takes six months of work.

Expected value

If the intervention has extremely high expected value, then the risk will be more favorable. We can compromise on evidence because the potential payoff would be astronomical.

Tricks people play with evidence

Evidence, like most other systems, can be gamed and made misleading. **Entire books have been written** on how to lie with statistics and how to analyze studies. These are worth reading but in the interim here are a few fast tips that will save you a world of time when analyzing data.

1. **Be wary of popular science summaries of research.** Popular science has a different goal than normal science: its goal is to be read more than to be accurate. This tends to lead popular science to make much stronger claims than what a proper systematic evaluation would show. A study claiming that “Chocolate is healthy for you” is far more interesting than one that says “Some minor negative effects and some minor positive effects found in mice eating chocolate.” Whenever something starts to seem a bit too polished or shocking it often ends up to be stronger claims than the original science. This sort of pop science is also extremely common on charities websites.
2. **Be wary of abstracts.** An abstract is a handy tool for quickly understanding a study, but much like pop science it can often overstate the claim. For studies that are really important to your charity's impact you will have to dig a bit deeper.
3. **Check effect sizes.** “Statistically significant” is what many studies focus on, but unfortunately this number is just not enough to make important claims. An effect size shows the magnitude or importance of the effect. For example, it could be statistically true that eating a certain food reduces cancer risk. However, if the effect size is extremely small, it would likely not be worth changing your diet. With respect to charitable interventions, it might be true that clean water reduces diseases, but the really important question is: by how much does it reduce them – i.e., what is the effect size?
4. **Do a multiple analysis correction:** A huge mistake a ton of researchers make is to measure and analyze a large number of effects without statistically taking this into account. A simple example of how this could go wrong: if you are measuring the effect of jelly beans on cancer and you measure and analyze a single variable, you will find jelly beans have no effect. However, if you break the jelly beans down into groups by color and test each color individually, you will have a much larger chance of finding an effect by pure chance even if no effect exists. There are very simple ways to fix this problem, such as using a Bonferroni correction. If a study measures a lot of factors it's worth using something like this to see if the effect holds up.
5. **Do a multiple study correction:** The same problem described above can also happen with multiple studies. If you see a single study that says a claim for example that chocolate is healthy for you make sure to check the other studies to make sure that there are not 100 studies saying it's bad for you and 1 saying it's good (in those sort of cases it's very likely the single study is wrong not the 100 previous ones).
6. **Replicability:** In general be careful with how often you expect a study to replicate. Many people assume that if it's statistically significant at 95%, it will replicate 95% of the time in future studies. This is very far from true. In

fact, a p-value of 0.04 means that it'll only replicate on average 25% of the time. This is part of the reason why multiple studies converging on the same answer is so important.

7. **Regression to the mean:** Effects generally get weaker as they get studied more; this is a universal statistical phenomenon. Extreme initial measurements tend to become closer to average the next time they are studied. This should make you very cautious of single pieces of evidence that show unbelievable effects.

V. Tools of effective altruism (the wrench)

A relative newcomer in the world of decision-making, the effective altruism movement combines tools from economics and philosophy, often aimed at applying your impact and causing the most good with the smallest amount of resources. Like a wrench, many of the techniques of effective altruism are about amplification and the maximum effect.

The techniques of effective altruism are already uniquely built around doing good and charitable output, so a large number of them are directly applicable to founding a new effective charity. In this section, we will cover:

- Counterfactual impact: what would happen if you didn't act?
- The paradox of seeking option value
- How to apply equal rigor of methodology to effective interventions
- How to consider which limiting factor will stop your charity from achieving scale

1. Counterfactual Impact: What Would Happen If You Didn't Act?

Author:

CE Team

Counterfactual impact is what would happen if you did X compared to what would happen if you didn't do X. For instance, if you brought chlorine to a village in rural Kenya with no safe water source, you would probably feel like you'd accomplished some good in the world. If, on the other hand, you gave the same gift to the average Swedish city, they'd probably give you a funny look – They already have chlorinated water.

Let's take a look at the counterfactual scenario (i.e. what would have happened if you hadn't donated the chlorine?). The Kenyan villagers would have continued to drink contaminated water and die from waterborne diseases. The Swedes, however, would have continued to drink their clean water. In fact, even if a Swede drank

straight from a contaminated river, they could likely go to hospital, get the required treatment, and survive to tell the tale.

This article will explain how to make sure that your counterfactual impact is as high as possible.

Neglectedness

When evaluating your counterfactual impact as a charity entrepreneur, the biggest factor to consider is other charities in the field. The extent to which a cause area is ‘overlooked’ is often called ‘neglectedness’. These are the three big questions you should be asking yourself:

1. Are other charities implementing this intervention?
2. How prevalent is the problem?
3. Why are there gaps in the first place?

Are other charities implementing this intervention?

Each year, innumerable well-meaning people start new charities in spaces that are already crowded. One of the first things to do when considering an intervention is to check if someone else is already implementing it in that location and in an effective way.

Be careful not to overestimate the size of the other organizations in the space.

Nonprofits often put very optimistic numbers for coverage rates. For example, if they say they have one million people in their program, it might be safe to assume they have efficiently and fully covered 100,000 people. Any claims that are hedged with “we expect” or “we plan on” should be taken with even more skepticism.

Another thing to bear in mind is **even though charities may exist in the field, many geographies may not be covered**. If one nonprofit is promoting fortified flour in Kenya, this doesn’t mean there isn’t room for another doing the same in India. In fact, even if there is already one in India, the size of the country means there could be states bigger than most nations that are still underserved.

How prevalent is the problem?

Even if someone is already working on the same intervention in the same place, that doesn't necessarily mean they are successfully fulfilling the needs of the entire target population. An orphanage in the Ukraine does not guarantee that all Ukrainian orphans are covered. It is important to factor in the prevalence of the issue in addition to the number of charities currently in the field, so you can work out whether there's an unfilled need. It's very common in the charity sector to forget just how large the planet is, and lose heart if even just one charity out there is doing something similar to you.

Bear in mind that when we talk about **neglectedness**, we mean **in comparison to other cause areas**. There are many unresolved problems in the world, but some receive attention disproportionate to their severity. For example, cancer is a huge problem and one could say it is neglected in the sense that it is unresolved, but under our use of the term, it would not be categorized as a neglected cause area. This is because compared to similarly pressing issues, there are more people working on resolving cancer. So the implication is that a comparatively less well-known or populated cause area has more **low-hanging fruit** and, therefore, investing X amount of capital (be it time, money, or effort) in solving a more neglected problem will go further than investing the same amount in a more mainstream one.

Why are there gaps in the first place?

If you've researched the other charities currently operating in the field and identified a gap you could fill, the next thing you should do is find out why that gap exists. Some potential reasons include:

- It's a new idea, so it's taking a while to 'catch on'.
- It requires a particular value set or epistemic stance that is rare.
- The organizations that would like to fill the gap don't have the funds to do so.
- The unserved areas are difficult to help (e.g. cultural barriers, conflict zones, lack of infrastructure, etc.).
- An organization has actually already committed to filling the gap – it's just a matter of weeks or months until they do so.

If the gaps are there for reasons that you can overcome or reasons that do not apply to you, great! If you uncover a persuasive reason why you shouldn't go into the field, great! You have just saved yourself months or years of time.

Often, you won't find the answers to these sorts of questions online, so make sure to **reach out to experts working in the field and ask them directly**. You might assume they wouldn't want to talk to you and give away 'trade secrets', but most charity workers are, well, nice people. They have the same goals as you and more often than not, they will be willing to answer any questions.

Speeding things up

The next thing to take into account is that if you come up with a really excellent intervention, people will eventually start copying it. This may be a very good thing for the world, but it does make it trickier to establish your counterfactual impact.

As an example, consider the 19th century debate about who discovered the theory of evolution, Darwin or Wallace. Most likely, despite being separated by thousands of kilometers, both brilliant men came to the same conclusion at roughly the same time. If Darwin hadn't been the first to publish his theory, Wallace would have done so shortly afterwards. So, even for something as unintuitive and revolutionary as the idea of evolution, Darwin only sped up humanity's knowledge by a few years.

All else being equal, your counterfactual impact will usually be much higher if you focus on neglected cause areas and problems that are very far from solved.

Generally speaking, **the more unusual the values or epistemics your intervention requires, the longer it will take society to 'catch on'**. For example, if you were to get in a time machine and travel back to the 1500s to promote the abolition of slavery, you might be able to speed up the process by centuries. You also might be ignored, ostracized, or worse, which is one of the risks associated with being ahead of your time.

A note on motivation

You might be feeling discouraged after reading all about how someone someday will probably have the same brilliant idea as you. Try not to focus on the forest and instead zoom in to the individual trees. If you say to yourself, "I'm only moving this field forward by ten years", that doesn't sound so good. However, **zoom in and look at the individuals you are helping**. Think of the mother whose child has just died of a vaccine-preventable disease. She's hurting right now, and telling her that others in ten years won't suffer the same does little to comfort her. You can also think of your own life. If somebody saved your life, would you dismiss it as, "just helping one person", or would you be immensely grateful?

Behind statistics are real lives. If you move a field forward ten years in expectation, you are likely making a huge positive impact on the lives of many individuals.

2. The Option Value Paradox: How Option Value Can Ruin Your High-impact Career If You Have Not Considered It Carefully

Author:

Joey Savoie,
Co-founder & Director of Strategy

A key factor to consider when choosing a career is option value – the concept of “keeping doors open”, just in case you would like to take advantage of alternatives in the future. However, a trend I have noticed in the effective altruism movement is that people will turn down other options due to the fear that they will no longer be able to return to their current job – a job that, ironically, was in turn chosen in part due to having the best theoretical option value. In other words, **the theoretical option value of the job is different from the practical option value.**

As an illustration, imagine two jobs – Job Always Open and Job Small Window. Job Always Open is flexible and allows you to enter it at any point. In contrast, you must enter Job Small Window within one year of graduating college, and if you ever leave the job, you can never get it back. The window of opportunity to enter it is small. It looks clear that all else being equal, Job Small Window gives more option value because it keeps open both options. This gives it the highest theoretical option value.

However, the equation starts to look much less clear once human psychology, more than one job, and multiple iterations are taken into consideration. In the real world, it's not the theoretical option value that matters: it's the ability to move flexibly between jobs or move to higher-impact options when they appear, resulting in the most expected utility over the long run.

If you choose Job Small Window, you later face a new dilemma. Job New Opportunity has come up and it's far more impactful than Job Small Window. However, you have invested time in Job Small Window, and you understand that if you leave it, you cannot return. Job New Opportunity is more flexible, and jobs like it will come up once every five years or so. However, you will be starting at the bottom of Job New

Opportunity. This means that, during those five years between opportunities, you are gaining career capital in Job Small Window. As a result, the next time a job opening for Job New Opportunity opens, the difference in appeal between the two jobs is even larger.

In practice, Job Small Window ends up leading to dramatically fewer options. Taking into account fear, sunk costs, and constantly weighting option value highly, Job Small Window might win every job comparison, even with jobs that would score far better in both the long and short term. **Consequently, jobs taken initially for option value can often reduce the number of options someone feels comfortable pursuing.** I have seen this most commonly in academia, finance, and technical careers, but I think it can happen with any non-flexible field or job.

To give a concrete example: Jane is an intelligent effective altruist who wants to do good in the world. She considers a lot of factors, but it's unclear which gives the highest impact. She decides she will go into a competitive career in finance because she figures that it will be easier to move from finance to other options than vice versa. This option value weighs heavily in the calculation, ending up as one of the primary reasons she chooses this path. However, as Jane gains more experience in the field, she feels that to keep this finance path open, she will have to focus heavily on that path, turning down other more impactful options and limiting the energy she can put into even looking for these options.

Each time another option is presented, even if it looks more promising, Jane turns it down because she thinks that once she leaves finance it will be very difficult to reenter it, but the other options are always available to enter. She would only take another job if she decided to leave finance forever. However, this is a high bar to reach both psychologically and practically; impact calculations are messy, and each year Jane feels more committed to her career path.

Sadly, this equation does not change over time. Quite the opposite: as she progresses, Jane feels increasingly invested in this career and less and less able to move from it. Forty years later, Jane looks back, having turned down dozens of far more impactful jobs than the one she ended up doing. Each time, she wanted to maintain theoretical option value. However, **now she realizes that her first choice limited the practical option value.** Her psychology, chosen career path, and calculations closed doors over time instead of opening them.

Option value is not a goal in itself; it is merely a means to an end. For effective altruists, that end is a higher-impact career path. However, if a job promises option

value precisely because it is hard to leave psychologically or logistically, this can make it very hard to pursue those paths. The job can give theoretical rather than practical option value. **Being able to prepare for, consider, and follow higher-impact options is the option value that really matters.** The theoretical option value may, in some cases, correlate with this, but in other, paradoxical cases it might be inverted.

The solution?

There are a few techniques you can use to reduce the odds of getting stuck in an **option value paradox**. One is being aware of the paradox and aiming to maximize the practical, rather than theoretical, option value. Another is trying to map out early what forms of reasoning will continue to apply across your career, compared to those which will change if and when you find a more impactful career path. For example, the option value conferred by a degree is different from that of a career in finance, as a degree has a clear endpoint.

To reduce the effect of the sunk cost fallacy, a good question to consider is, “If I weren’t in this career right now, would I choose to enter it again instead of choosing this other option I’m thinking of?” You can also ask external experts, particularly ones with shared values, about comparing the two career paths. This can reduce the effects of personal anxiety and sunk cost. A final strategy is to watch out for careers that give good option value but are very hard to leave, either psychologically or logistically. Option value is important, but it can also ruin an impactful career if you do not consider things carefully.

3. Equal Application of Rigor for EA Interventions

Author:

Ishaan Guptasarma,
Curriculum Developer & Content

Slatestarcodex has an interesting post regarding the [equal application of rigor](#) in philosophy. This principle can be applied in a practical sense to effective altruist interventions as well.

The case of the life-saving bus

A great example of difference in rigor is the different ways charities analyze data. There is a common bus ad that says “Save a life for \$10” with a picture of a child presumably similar to a child that could be saved. Now this is a pretty exciting claim if taken at face value. GiveWell top charities currently save a life for \$3,500, so this charity on the bus in theory could be many times more effective. But of course the difference here is not in impact but in measurement methodologies and rigor. The charity, it turns out, gives out a life-saving bednet that could save a life, but if carefully evaluated many many nets need to be given out per life saved. And in fact GiveWell’s top charity AMF gives out nets for about \$3 instead of \$10. This makes the bus charity less than half as effective.

The difference is that GiveWell takes into account a much wider range of factors, including how many nets need to be given out on average to save a life. If you took the raw numbers, saving a life for \$10 seems like a far better deal than saving a life for \$3,500. But on closer inspection, giving out a bednet for \$3 seems better than the same net for \$10. Many of the differences you see in the charity estimates and impact evaluations come from differences of methodologies not differences in the charity. This is part of why we should not get excited every time we see weak evidence (like a bus ad) of a new charity that is highly effective. It’s most likely a weaker methodology rather than a hidden gem of effectiveness.

The case of mosquito nets

Imagine that an established charity evaluator estimates the degree to which mosquito nets protect against malaria, estimates the degree to which a net distribution scheme leads to the actual use of nets among the population, estimates the prevalence and lethality of malaria in the region, combines these numbers to estimate the amount of malaria prevented and the number of lives saved, and thereby comes up with a cost per life saved estimate for mosquito net distribution based on those numbers. Is this number accurate?

You can't be certain. Even if all of the above estimates are accurate, the true amount of good that is done might be greater than that. Not having malaria could improve developmental outcomes for children by a quantifiable amount, grant parents greater capacity to earn income and take care of children, shield people from the mental health trauma of losing loved ones in a way that makes a measurable life satisfaction difference, and so on. Your original analysis may not have accounted for this, and if you add it in, your number might change.

The true amount of good that is done might also be less than that. You might discover that some fraction of the nets are used for less effective activities like making fishing nets or clothing, or end up being sold on a market where they do less good, decreasing the efficiency. You might discover that helping people earn more raises prices and hurts the neighboring villages, resulting in a small negative effect. Your original analysis may not have accounted for this, and if you add it in, your number might change.

The freedom to choose which secondary effects you choose to account for allows you a lot of control over the end result of a cost-effectiveness analysis, regardless of how thorough and accurate the analysis is.

The more positive secondary/side/flow-through effects you add to your analysis, the more good it will seem like you're doing. The more negative effects you add, the less good it will seem like you're doing.

Opt for equal comparisons

You might think that when choosing an intervention, you should account for as many positive and negative outcomes as possible, so as to get the most accurate

number. While it is good to practice this sort of thoroughness for the sake of improving your knowledge and identifying any crucial considerations which may dramatically influence your effectiveness, when making the final comparison between interventions, **you want to avoid a situation where the numbers you get are primarily an artifact of the methods you used to get them.**

Therefore, when making comparisons between interventions you should try to make sure that your analyses are as parallel as possible. **You might not ever be able to say literally how many lives are saved per dollar, but you might be able to say that one intervention was three times as good as another** when the two interventions were evaluated side by side using as similar a method as possible. This doesn't solve everything of course – if you switch to a different method, even one applied uniformly across both interventions, 3x better might rise to 10x better, or drop to only 0.5x as impactful. Applying equal rigor will unfortunately not remove all forms of model uncertainty. However, it is a good tool to keep in your arsenal and it will help prevent one form of inaccuracy.

Watch out when assessing your own ideas

One area where it is especially important to remember to apply equal rigor is when evaluating your own work and your own ideas. When you are directly working on something, you can often see with your own eyes the positive impact that it has and observe some extremely good flow-through effects that may not be captured by your cost-effectiveness analysis. The temptation may be then to adjust your analysis to take the evidence of your eyes into account, come to a high number of cost-effectiveness, and trick yourself into thinking that your intervention is definitely the best use of resources on the margin. However, what's really happening is that you see all the good that your own project is doing in very high definition while you look at other projects in terms of only summary statistics that don't capture everything.

If you find yourself comparing your own projects and ideas against others and notice that yours seem unusually good, **consider if you're quantifying the good aspects of your own projects with greater rigor due to being able to see them better.**

A person who gives money to impoverished individuals who they personally know can easily list all the positive effects on each and every one of them: "I gave \$2,000 to John, and therefore he was able to rent an apartment and get his family medical care, and one of his children had a life-threatening disease which was caught early

thanks to this medical care, so, therefore, my cost-per-life-saved is at least \$2,000", may be an accurate statement. However, this is not a fair comparison because this individual can see everything that happens to John, whereas if for example a parent avoids malaria and, therefore, has more parenting capacity to get her kids medical care which saves their lives, you wouldn't see that particular life saved reflected on a GiveWell spreadsheet.

Of course, interventions are often very diverse and evaluating them using the same methods is hard. If they are somewhat similar, you can try to make them comparable (e.g. both vaccinations and mosquito nets protect against disease by some quantifiable amount, and share many similar considerations such as base rate of the disease in the population, lethality, etc). Another way might be to judge all interventions according to their one or two biggest and most straightforward effects (e.g. lives saved) and ignore all smaller indirect effects.

None of this is to suggest that indirect effects don't matter or should be ignored. You should try to quantify as many effects as seem useful to your decision. **The important thing is to make sure that comparisons between interventions are made with model equivalency in mind.**

4. Why We Look at the Limiting Factor Instead of the Problem Scale

Author:


Joey Savoie,
Co-founder & Director of Strategy

Scale, or importance, is held as one of the **three criteria** to consider when evaluating an intervention for promise. The idea is that large-scale problems might suggest an area will be more effective to work on (assuming it also scores well on the other criteria). Some interventions are predicated on very strong scale arguments, such as far future or wild animal suffering. However, we at Charity Entrepreneurship have found that scale specifically is quite a **poor indicator** of the promise of an area.


Organizations tend to be incentivized to scale both from an impact perspective and from a personal perspective. However, all organizations, eventually, hit a limiting factor that makes it hard to scale faster. For some organizations, it might be the total scale of the issue. For example, perhaps when working on a nearly eradicated disease, the limiting factor might just be how much more of the problem there is left to deal with. This would be a case when the traditional use of scale comes to almost the same result as the limiting factor model. Both models suggest this would not be a great area to work on due to the problem scale being quite small (currently).

However, there are many times when it diverges. For example, certain issues might have a massive problem scale, but be quickly limited by some other factor. GiveWell has talked about **surgeries being limited by the supply of surgeons** – this is not a scale of problem issue specifically, but it is a limiting factor. A scale model might suggest that if there is a ton of surgeries still to be done then this is a worthwhile issue to focus on. But a limiting factor model would suggest that it would be quickly capped by the number of surgeons. Overleaf is a very simplified comparison.

Scale model



Limiting factor model




As you can see, the results end up being quite different: surgeries are limited by the talent pool far earlier than vaccinations become limited by anything else. The magnitudes of the categories were set to be more cross-comparable (e.g. 1 million compared to 10 full-time staff). And both get limited several times before their problem size limit. In this case, it would not even matter if surgeries had 10 times the scale (say, 200 million people affected) if they will be stopped by talent, logistics, and funding before they can even help 20 million people. These numbers are estimates, but given our work and research in these areas we are confident both of these will hit a limiting factor far below their problem size (or what is often referred to as ‘Scale or importance’).

A common response is that scale is only one of the factors considered when evaluating an intervention. However, things like ‘logistical limit’ fall under the radar of tractability/solvability, and there are two problems with this. First, tractability is not really currently used in this way. Right now, lots of claims are being made along the lines of “cause X should be focused on more due to it having a huge problem size” with no further reference to tractability. Secondly, tractability is often seen as the speed at which you will make progress as opposed to a specific factor that will stop growth from occurring. For example, an intervention could be very shovel ready, but only at a small scale before its limiting factor comes into play. If we go back to the surgery charity example, it could be that 3 surgeons want to start a charity, and for them it is very tractable and shovel ready, since their marginal effort is great and, at a very limited scale, their solvability rate is very high. The endeavour will run into specific “scale” issues involving a limiting factor of ‘hiring other surgeons’, and this ends up feeling more like a scale issue rather than a tractability one.

Another common concern might be funding limits, which on their own are not hard limits like problem size. However, while something like a funding limit can be improved upon with more fundraising and field building, this is not an easy task. For the money ranking, I do not think you would just put down the amount you have fundraised for the area, but a reasonable bound for how much could be fundraised, while taking into account the current donor space and a reasonable amount of time (e.g. 2-5 years). This can change over time, but so can the size of the problem: factory farming is a

growing problem and global poverty is a shrinking one, but that does not change the importance of having a sense of their scale.

A claim that I hear a lot in the animal space is that wild animal suffering is such a huge scale problem that we should seriously consider working on it. Many people would suggest that wild animal suffering is an issue with a much larger scale than something like vaccinations based on the pure number of beings affected. This claim is definitely true, since there are trillions of wild animals and only ~20 million people in need of any single vaccination. But if we look more closely at its limiting factors, I think this claim is pretty misleading.



The problem size limit is indeed huge. In fact, it has to be cropped, or otherwise it would go all the way to the top of this article and the other sections would become impossible to see. However, that is not really what matters. Even if wild animal suffering is a huge problem and even if there is only a very limited amount of funding and talent that wants to work in the area, you will bump into problems with scaling far before it starts to matter if there are a billion wild animals or a billion billion. In the practical cases, when comparing these interventions, if a charity were founded in both of

these areas, the vaccination charity would be able to get to a much larger scale than the wild animal suffering-focused charity could. A claim that “we should work on intervention X due to its massive problem scale” seems quite inaccurate. These sorts of arguments are extremely common for wild animals specifically and more broadly in EA.

Research to determine a cause’s limiting factors generally ends up spending more time on considering the funder and talent space compared to, for example, mapping out the specific number of animals affected by any given intervention. And it usually ends up with a set of fairly different interventions looking promising.

Of course, the perspective changes depending on what you are looking to do in a given area. For example, when considering donating to a charity, the main thing examined by GiveWell is room for funding. They are looking for an intervention for which its limiting factor to doing more good is funding. Some interventions might be very promising, but due to talent, logistical problems, or the size of the problem limitations, even if GiveWell gave them more funding, they would not necessarily be able to create more impact. Organizations with room for more funding are generally stopped not by other factors but by ‘room for funding’ itself. Due to GiveWell being a funder, it makes sense for them to take a careful look at room for funding as that is the limiting factor they can improve. As a charity entrepreneur, what the funding space looks like for a given organization should play a big role in what charity to found. As an employee at an NGO, you mostly have to consider how much of a limiting factor talent is for the organization.

These models are really simplified on both the scale and the limiting factor side, and I think there are other possible ways to use scale differently (or use tractability in a different way to cover some of the same concerns). I do not think 100% of effective altruists use a simple scale-based way of looking at problems, but I do think a large percentage of effective altruists use a fairly simple “size of the problem” based way of considering scale without thought to the percentage of the problem they can solve, given the first limiting factor that will stop growth/progress.

VI. Long-term planning (the pencil)

Every great project needs to be sketched out before its creation. Thanks to a pencil, you can draw and draft but erase and change as you learn more. The same will be true of your long-term plan, which will require updates as you go. All great plans will have parts that need to be erased and iterated many times at many different points throughout the year. However, having a drawing of what your endline project will look like is equally crucial both for decision-making and practical projects. In this section, we will cover:

- The importance of a strong theory of change
- How to make different levels of plans, such as 5-year, 1-year, and month to month
- How to consider going for-profit or not in the long term
- How large to scale, or if to scale up at all

1. Theory of Change

Author:

Ishaan Guptasarma,
Curriculum Developer & Content Writer

A “theory of change” is a document that explicitly articulates the assumptions that underlie your plan to take certain actions toward achieving a specific goal, and lays out a method to test them. It is generally represented in the form of a cause and effect diagram. A well-designed theory of change allows you to clearly communicate what your activities are, and why they lead to the outcomes that you and your supporters want.

Example 1: Research organizations

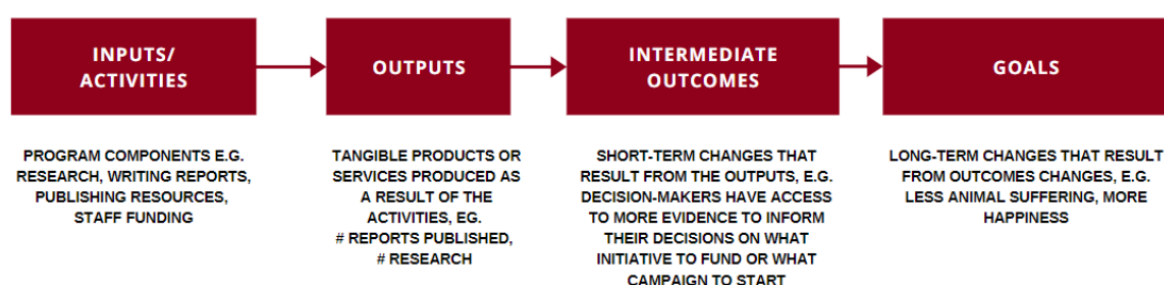
Imagine that you wanted to start an organization that uses research to make a difference. Without an explicit theory of change, your implicit mental model of how you might make an impact may go something like this:

Do reviews of literature → Find insights that seem useful → Publish them

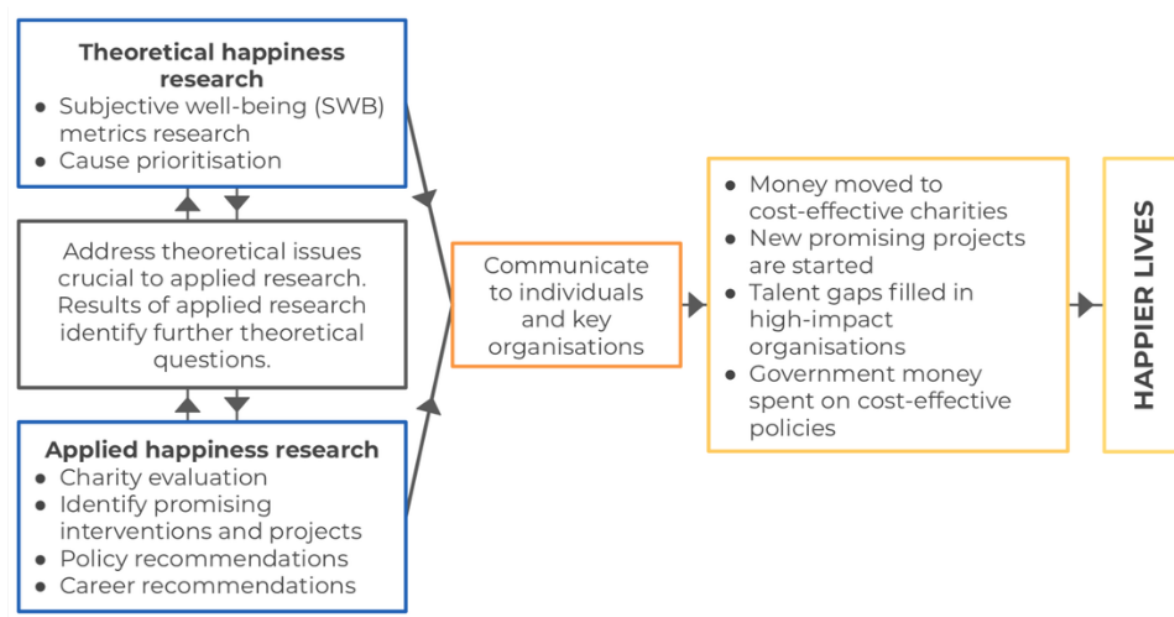
This is not a good theory of change. It doesn't properly outline your goal, it doesn't explain how your actions will lead to that goal, and it doesn't explain how you will measure what you are doing. Presumably, your final goal is to positively impact living beings, not just publish papers. Even if you work very hard and put out many widely cited papers, it's possible that you will not make any difference.

Check out this EA forum post by Karolina of Charity Entrepreneurship about [increasing the impact of your research](#) to see how this theory of change could be improved.

A good theory of change draws the full causal chain from your actions to the final impact, which is your end goal.



After going through Charity Entrepreneurship's curriculum, the Happier Lives Institute published a [theory of change](#) that follows the above principles.



On the left, you can see the core research agenda of the Happier Lives Institute at a glance (the inputs/activities). On the right, you can see some measurable metrics such as money moved, new projects, and hires made (intermediate outcomes) which lead to the final impact – happier lives.

It is clear from this diagram that “communicating to individuals and key organizations” is an important part of this process. Ultimately, the goal of doing research is so that *someone* will read it and become able to allocate a grant in a better way, or start a new project in a more promising area.

Making this explicit is important because it allows the research agenda to proceed with clear goals and objectives in mind. If you are well aware from the start that one of the major goals of a specific research project is to improve grant allocation, you can contact those grantmakers before beginning the research and get a sense of what sort of information might help them make better decisions.

If you do this process, your research will be much more likely to have an impact than if you simply publish a research project covering a topic in the field that seems interesting and hope that someone will one day find it useful.

Example 2: Cash transfer organizations

Here is a theory of change for an organization that uses cash transfers to increase immunizations. They have two “input” actions: providing resources to keep immunization camps reliably open, and providing cash incentives to people who get vaccinated.

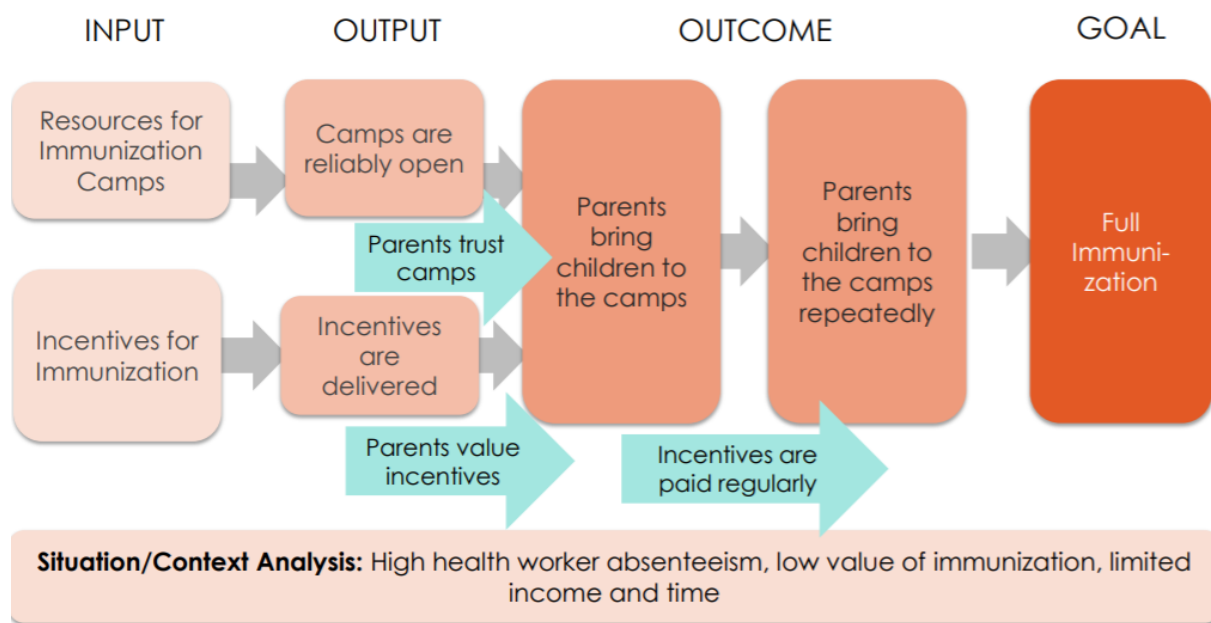


Image source: Aparna Krishnan (July 3, 2017) *Theory of Change* [PowerPoint lecture at JPAL South Asia at IFMR.]

There are also several *measurable* “outputs and outcomes”. You can send someone to check and see if a camp is reliably open. You can create a verification system to check if incentives are delivered. You can count how many children come to the camps, and how often they return for their booster shots. If something is going wrong, you can check your assumptions (e.g. surveying parents to see if they trust the camps).

When designing measurement systems, keep in mind how much they cost. If a measurement is expensive but isn’t meaningfully informing your decision on whether it’s worth investing more resources in this intervention, cut it. Do try to measure things in more than one way, however, especially if there’s scope for a particular measurement to be misleading. You can learn more about this in our [monitoring and evaluation](#) article.

Convert to Log Frame

	Objectives Hierarchy	Indicators	Sources of Verification	Assumptions / Threats	
Impact (Goal/ Overall objective)	Increased immunization	Immunization rate (% of children immunized)	Household survey	Adequate vaccine supply, parents do not have second thoughts	<div>Needs assessment</div> <div>↑</div> <div>Impact evaluation</div> <div>↑</div> <div>Process evaluation</div> <div>↓</div>
Outcome (Project Objective)	Parents attend the immunization camps repeatedly	% of beneficiaries attending follow up camps	Household survey; immunization card / attendance	Parents have the time to come	
Outputs	Immunization camps are reliably open; Incentives are delivered	% camps open as scheduled; # of kg bags delivered;	Camp administrative data; Random audits	Nurses/assistants will show up to camp and give out incentives properly	
Inputs (Activities)	Camps + incentives are established	% of Camps built, % of camps functional	Random audits of camps	Sufficient materials, funding, manpower	

Needs assessment 43

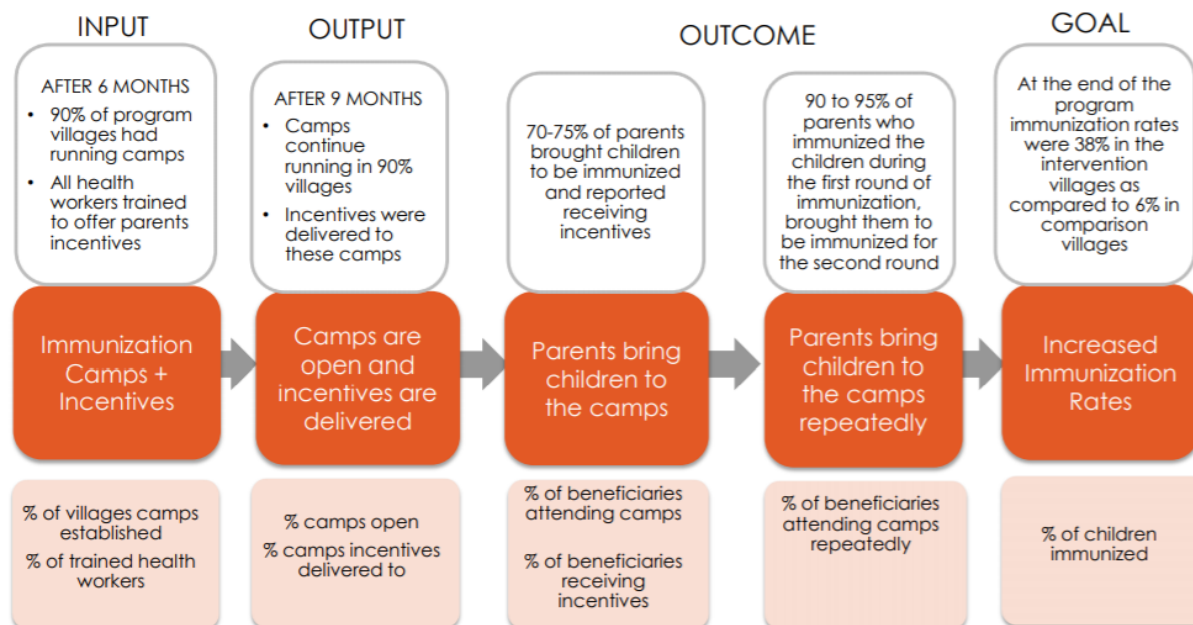


Image source: Aparna Krishnan (July 3 2017) *Theory of Change* [PowerPoint lecture at JPAL South Asia at IFMR.]. See also our [logframe template](#) discussed in our [M&E](#) section.

Common errors to watch out for

Don't forget the endline metric. The above example is a good theory of change, but can it be improved? Immunization isn't the end goal – saving lives and improving quality of life is. This may *seem* obvious, but it is very important to say it explicitly, so that you remember to model how much disease is actually prevented and how many lives are actually saved by increasing immunization in a region.

If you skip this step, you might (for instance) end up choosing a location that maximizes the number of people you can immunize, rather than choosing a location that maximizes the number of lives you can save (e.g. by picking a location with a high child mortality rate). Effective organizations are built by systematically examining many simple, perhaps obvious-seeming considerations, and the theory of change is an important step in that process.

Beware of mission creep: This theory of change actually has two interventions: opening camps *and* incentivizing people to attend them. Sometimes, this is appropriate. But suppose it turns out that building camps provides most of the benefit, even without the incentive? In that scenario, you might divert the money that goes into incentives and use it to build more camps, and vaccinate more people and save more lives. Or, suppose it turns out that incentivizing preexisting camps would provide most of the benefit of creating new camps at a fraction of the price? In that scenario, you could divert resources you might spend building camps toward just handing out conditional cash transfers, and therefore vaccinate a lot more people and save more lives.

Suppose as this hypothetical nonprofit grows larger, they begin to think: “If we have the camps, why not also provide vitamin supplements, or distribute essential medicines, or have the workers build latrines to prevent the spread of disease?” Sometimes adding more programs is efficient, squeezes more impact out of fewer resources, and does make sense! It also often looks good on paper and impresses some donors to have a more comprehensive and holistic program. However, be sure to do an explicit cost-effective analysis first. There are exceptions, but *generally* speaking, if you are considering multiple interventions, one of those interventions will be more efficient than all the others and you shouldn't divert any resources from it.

Stay specific and skeptical: Don't have any hand-wavy steps in your theory of change. Some examples of hand-wavy steps are:

Publishing research → lives are saved. As covered above, you need to add steps where you explicitly model who is going to use it and what actions they will take, and you need to actually talk to these decision makers. And, you need to quantify the impact of these decisions on living beings.

Raising awareness → less suffering. You should explicitly model and measure or at least estimate the rate at which this translates to behavioral changes, voting changes, policy changes, etc. And, you need to quantify the impact of those changes on living beings.

Changing people's views → impact. You should outline the mechanism of how changing attitudes influences living beings, and estimate or measure the degree of this impact.

Don't just assume that one step will lead to another. Don't overestimate the probability that one step will lead to another. Find an objective way to verify it. Get feedback from someone you know to be skeptical.

Resources

GiveWell: [Why it's important to think through all of the factors that influence a charity's impact](#)

[Are you working on a research agenda? A guide to increasing the impact of your research by involving decision-makers](#)

[HLI has hatched: strategy update after the Charity Entrepreneurship Incubation Program](#)

Mary Kay Gugerty et al (2016). [Guiding Your Program to Build a Theory of Change](#). Innovation for Poverty Action

Harries, Hodgson, and Noble (2014) [Creating your theory of change: NPC's practical guide](#). Think NPC

Aparna Krishnan (July 3, 2017) [Theory of Change](#) [PowerPoint lecture at JPAL South Asia at IFMR.]

Research to Action: [Introduction to Theory of Change](#) (video)

2. Long-term Planning

Author:

Ishaan Guptasarma,
Curriculum Developer & Content Writer

As an individual and as an organization, you should have long-term goals about what you want to accomplish. What changes do you wish to bring to the world this year? What about in five years?

Your **theory of change** specifies certain actions that you will have to do in order to accomplish your organization's goals and achieve impact. Many of these actions may involve complex projects that require doing a sequence of smaller tasks over a long period of time. However, most people have a tendency to underestimate how long a task will take and underestimate difficulties. This phenomenon goes by many names ("planning fallacy", "Murphy's law") and it can lead you to go off track and fail to achieve your goals due to pursuing projects you don't have time for.

Making a plan will help keep you organized and focused on your highest priority tasks (see a framework for identifying high-impact tasks [here](#)). Even if you work a lot, it's easy to get sidetracked with stuff that is not important for your main plan. Keeping track of time will help you use it well. By making estimates and keeping track of how long you thought things would take vs. how long things actually take, you will improve your ability to strategize about actions taken by your organization. Plans can also encourage you to anticipate problems before they happen. We recommend doing "pre-mortems" as part of your planning: imagine that your charity failed, think of the most likely reasons why this might have happened, and come up with some ways to prevent this outcome.

We recommend three levels of planning – an aspirational 5-year plan, a brief but precise 1-year plan, and a more detailed month-to-month plan.

Five-year plan

Here is an example of CE's **5-year plan**.

You will notice that we have created specific benchmarks for what we want to accomplish, as well as how many staff it will require (and therefore the approximate cost). By stating your long-term goals and criteria for success and failure beforehand, you will keep yourself honest about whether or not you are meeting your goals. In this way, you can avoid “moving the goalposts” – don’t keep changing the criteria for “success” to fit whatever you did achieve, set them beforehand. For instance, if your **counterfactual impact evaluation** specified that you had to achieve a certain goal in order to consider this activity sufficiently high impact to be worth doing given the alternatives, and 5 years later you did not achieve it, it will be easier for you to acknowledge this to yourself if you set the criteria beforehand.

You might also notice that we have a plan for testing and scalability. In addition to being directly important for realizing long-term impact, having a clear vision for the longer-term future is important with respect to communicating with donors and grantmakers. Many will want to know whether the project will be **sustainable** (will eventually be able to continue without their support) and **scalable**.

Things are guaranteed to change over time, so there's no need to be too detailed and specific with a 5-year plan. It should provide you with a solid strategic trajectory and should guide your more immediate goals, including in the one-year plan.

One-year plans

Here is an example of CE’s **1-year plan**.

A one-year plan tends to include the following elements:

- A set of very specific, measurable goals and clear timelines for achieving them. This plan also includes a plan for how progress on the goals will be measured.
- A **budget**. Funding typically is given out on a yearly basis, so 1-year plans are often part of a fundraising ask. Note that we have also included a rough outline of who we are hoping will fund this budget.
- An explanation of how these goals tie into past activities and connect with longer-term future goals

You will often need to make a one-year plan as part of a grant application. Keep it concise (1–3 pages) so that other people can read it easily. It helps to talk about how this year’s activities tie into the bigger picture over the next five years.

Month to month timeline

You should plan out roughly what you will do for individual months. This becomes especially important as your organization begins to have more staff members. Each staff member has a limited number of hours, so it helps to make sure that this time is being allotted to the stuff which will help them have the highest impact.

Below is a template of a **month-to-month timeline**. Each project should be planned out in terms of when it will begin and end. You might also specify who or how many full-time staff equivalents will be working on the project. Your timeline doesn't necessarily need to be very detailed in the early days, but it's good to have a sense of what everyone is working on during any given time period. You might also encourage individual staff to make weekly plans for themselves, and share it with those who collaborate closely with them.

Month to Month plan - Template			2019													
			Q4					Q1								
			DEC					JAN					FEB			
			29	6	13	20	27	3	10	17	24	31	7	14	21	28
Area	Activity															
1	Operations	Registration														
		Item 2														
2	Category B	Item 3														
		Item 4														
		Item 5														
	Category C	Item 6														
		Item 7														

In addition to helping you plan, having a legible timeline can help your staff understand the internal parameters that they are working with. The more information your staff has about the big picture, the more connected they will be to the work, and the better they will be able to take the initiative, manage themselves, and prioritize their time correctly.

Tracking progress and evaluating next steps

Every month, it's worth going through your monthly and 1-year goals. Try highlighting things that are "on track", "not on track", and that "seem unlikely to happen" in different colors. After seeing how the month went, you can refocus. You should also do this process before adding any new goals that are likely to take a lot of time, so you can remind yourself how long things take.

You should also build longer, more in-depth **reevaluation points** where you deeply question all your core assumptions. Which goals did you meet this past year? Do you still think this is the highest impact project you can work on, and if not, how can you pivot or scale down? Which aspects of your process worked well, and which need improvement? Did you learn any new information, and how should you change what you are doing based on that info? Reevaluations should also be done when a major organizational change occurs or a key new piece of information is discovered. Check out CE's **yearly review** for an example of this.

Summary

We recommend making a 5-year plan, a 1-year plan, and more detailed month-to-month plans. Some benefits of making long-term plans include

- Staying on track and not getting distracted
- Calibration and realistic estimates, avoiding the “planning fallacy”
- Anticipating problems (do “pre-mortems”)
- Being presentable to donors and grantmakers
- Commitment and accountability, avoiding “moving the goalposts”
- Keeping staff informed about the big picture agenda

We also recommend tracking progress and scheduling regular reevaluation points.

Resources

[CE 5-year plan](#)

[Monthly plan template](#)

[CE 1-year plan](#)

[CE 1 year update review 2019](#)

Handbook

[!\[\]\(0d7ca0919e6c47bbd874bfa0189fe22e_img.jpg\) Financial Planning and Budgeting](#)

[Theory of Change](#)

[When to Focus and When to Reevaluate](#)

[What is an Impact Evaluation?](#)

External

[SMART Goals](#)

[The Secret to a Great Planning Process — Lessons from Airbnb and Eventbrite](#)

[Set goals with OKRs \(Objectives and Key Results\) – Google](#)

[How Timeboxing Works and Why It Will Make You More Productive](#)

[The Founder's Guide to Discipline – FirstRound](#)

3. Going For-profit

Author:

Antonia Shann,
Editor & Communications Officer

Note: This article builds on major contributions from Peter Brietbart, CEO of Mind Ease; thank you.


You may know the **big names in EA** and in your cause area. You've sketched out some **early strategies** – maybe you've even drafted a couple of messages, or outlined a proposal.

Hopefully, the prospect of fundraising is starting to feel less daunting. But what if your organization provides value that people are willing to pay for? In this case, it's worth looking into starting a for-profit. After a quick look at different organizational structures, this article explores some of the pros and cons of going for-profit.

As with many binaries, the divisions between charitable and corporate organizations exist more along a spectrum:

- Some nonprofits generate a portion of their revenue through for-profit ventures. For example, some larger nonprofits earn money through charity shops, whether these be brick-and-mortar stores, or simply merchandise sold online (e.g. **Mercy for Animals**' shop). (Note that we don't expect this to be a common strategy among CE's incubatees, as it doesn't lend itself to early-stage organizations.)
- Some organizations are run on a purely for-profit basis, and are ultimately responsible to shareholders, not stakeholders. But they may nonetheless engage in a degree of corporate social responsibility (CSR). For example, take a look at **Microsoft's 2019 CSR report**. (Note that in India, **CSR is mandatory**.)
- A whole raft of organizations exist in the middle space. There are nonprofits that draw on private sector concepts (e.g. **Living Goods**) and for-profits with

a social purpose (e.g. **Impossible Foods**). There are nonprofits with for-profit subsidiaries (e.g. **Embrace Global**). And there are for-profit organizations that commit to an above-average level of social responsibility, defining success in terms of purpose as well as profit (e.g. **B Corp-certified** organizations like **Rising Academies**).

Where along this spectrum you ultimately situate your organization will depend on multiple factors. How confident are you that you can successfully align your incentives? What proportion of your funding might you get from for-profit ventures? What are the legal requirements and regulations in your country of incorporation? (Starting to feel lost? This is why it's so handy to get outside input, such as from a  **legal and advisory board**.)

In the majority of cases, **we at CE don't advise going down the for-profit route.**

- Many of the most cost-effective interventions suffer from market failure (e.g. vaccinations), and thus don't lend themselves to going for-profit. (Remember that **funding is usually not the limiting factor for EA organizations**.)
- Adding the goal of profit risks watering down your goal of impact, even if financial motives are only secondary. The issue fundamentally is that getting money through the door increases your chances of success, but not necessarily of improving the world.
- The **counterfactuals** seem weaker in a domain where profit can be made – money attracts talent, so your unique contribution is not essential. And although you might think that you could run a for-profit intervention more thoughtfully and impactfully than your counterfactual replacement, this point is both easy to rationalize and difficult to measure. Virtue signaling is a far cry from truly aligning profitability with impact, but it's a trap that many start-ups fall into.

Despite the concerns around going for-profit, it can be a good idea depending on your mission and its fit. A major advantage here is scalability. With your own resources, you don't need to rely on external funding (unless you go down the venture capital route). And if your intervention can fund itself, you remove yourself from the pool and avoid taking funding away from other effective nonprofits that suffer from market failure.

So in a few rare cases, going for-profit might be the move. If your intervention does fall into this category, what should you watch out for?

One of the major pitfalls when going for-profit is **misaligned incentives**. This problem arises when the way you make money and the way you make a difference are in tension. To mitigate this problem, it's vital to strategize on how to keep your business model in line with your impact. Your **theory of change** will help you out here, by ensuring you make explicit each step in your path to impact. Similarly, thinking carefully about your endgame (see our **next article**) will keep your goal in focus. Building transparency into your organization will also increase accountability in moving toward your goal.

(It's worth noting that misaligned incentives can affect nonprofits, too. The people you help are often not the people who fund you. There's a danger here that you overfocus on fundraising or on meeting the goals of your funders, to the detriment of your beneficiaries. This is why it's so important to have value-aligned funders – and, as with going for-profit, to reflect on your theory of change and your endgame.)

Going back to the for-profit/nonprofit spectrum, remember that there's wiggle room. Maybe you start out as a nonprofit but decide after a few years to move more in the for-profit direction. (Although getting into the **legal complexities of parents and subsidiaries**, for example, is probably taking on too much for the early days of your organization!)

Even if the for-profit route isn't for you, it's worth being aware of some of the concepts that might benefit your own work. Cost-effectiveness is, for obvious reasons, key among for-profit organizations – and it's also crucial for EA organizations.

- **Living Goods** is a good example of a nonprofit that draws inspiration from the for-profit sector. Their intervention pays community health workers to provide primary care, and helps them up their wages through selling health products. (Their **article on innovative finance** is also well worth a read.)

Ultimately, choosing between for-profit and nonprofit is about **which legal structure will best enable you to achieve your mission**. In most but not all cases this will be the nonprofit route. Keep a laser focus on your impact, and with a bit of elbow grease the rest will fall into place.

Potential for-profit ideas

Better chicken feed

A gratitude app

Biofortified crops

Contraception subscription in developing countries

Reading list

Internal resources

 Fundraising outside of Effective Altruism

The Case for Charity Entrepreneurship

 Prepare Yourself to Fundraise: Tips For Entrepreneurs

Theory of Change

 Layers of Fundraising and Alignment

External resources

What's Your Endgame?

Innovative finance

Should Your Business Be Nonprofit or For-Profit?

3 Reasons to Consider Converting a Nonprofit to a For-profit

4. To Scale, or Not to Scale: That Is the Question?

Author:

Antonia Shann,
Editor and Communications Officer

As we look for ways to maximize our impact, the enormity of the problems we confront yawns before us. Animal advocates are faced with the cold reality that seventy billion animals are farmed each year – not to mention the animals we pull from the oceans, whom we measure in tons rather than in individual lives. **One in four** of us will suffer from mental illness in our lifetime, and stigma prevents many from seeking help. Today, health care systems groan under the weight of a pandemic, as COVID-19 swallows up normal life.

Next to the statistics, our efforts can seem like a drop in the ocean. This is why scaling up is so urgent: the more we grow, the more we can help. But what if we're focusing on the wrong thing? In **this article** for the Stanford Social Innovation Review, Alice Gugelev and Andrew Stern shift the conversation away from scale-up and toward the concept of **endgame**. They observe that, while scaling up can certainly increase your organization's impact, scale and impact are imperfectly aligned. Essentially, through thinking consciously about our endgame, we can better map action onto impact, and do good more effectively.

Before exploring the concept of endgame, let's look at some of the challenges associated with scaling up. Stern and Gugelev point out that the barriers here are both immense and structural in nature. For a nonprofit, the leap from early stages (where budget is capped at roughly \$5 million) to breakout (\$5-\$10 million) or full scale (upward of \$10 million) presents an enormous challenge. The authors describe this gap as the **social capital chasm**, and highlight four aspects of the nonprofit sector that create it:

- Incentive structures (e.g. lack of equity/stock options) make attracting managerial talent difficult.

- There is usually no overlap between funders and direct beneficiaries, so charities must “win two games”.
- An emphasis on minimizing overhead can undermine operational capacities.
- Funding is erratic, since grants are normally allocated to specific programs rather than to broad missions.

As an EA organization, some of these challenges might affect you less, since staff are often motivated to accept lower salaries and overhead is less of an issue. But growth is often only possible outside the narrow world of EA donors. And even if an organization can overcome the social capital chasm, challenges remain. There’s scaling and there’s scaling well, for example. How do you ensure that you scale up sustainably and without watering down your impact?

At this point with snares on all sides, nonprofit work is starting to look a bit disheartening. So let’s turn to the endgame. First off, what does it mean to focus on the endgame? As Stern and Gugelev define it, the term endgame refers to a nonprofit organization’s intended role in solving a problem. While scaling up traditionally looks more at your organization’s direct impact, the endgame examines the part you play in shaping the broader sector. To illustrate the concept, Gugelev and Stern sketch out six possible endgames.


- **Open source.** An organization cultivates new ideas and interventions through research. Knowledge and resources are then shared with other organizations.
 - For example, Charity Entrepreneurship researches interventions through extensive research, and supports co-founders to start the most effective interventions through our Incubation Program. CE’s research and handbook for entrepreneurs are both publicly available to extend impact beyond program participants.
- **Replication.** An organization creates a model or product that can be easily reproduced. The original organization can offer certification and training, and act as a center of excellence.
 - Charter school networks in the US use replication centers to teach their model to other educators, whose preexisting infrastructure and embeddedness within a community means they may be better positioned to implement the model.
- **Government adoption.** This endgame is appropriate for an intervention that can be delivered at scale, and requires lobbying to influence policy and budget. After the intervention is adopted, the organization may continue in

an advisory role or as a service provider. Yet the ultimate goal is for the government to be in charge of financing and decision-making.

- **Charity Science Health** partners with state governments (as well as with other NGOs) to ensure the sustainability of their intervention, which involves sending SMS reminders for vaccinations. At some point, the government might be able to adopt the intervention itself, run it through the Ministry of Health and fund it with tax income. While the CSH model is very lightweight, government adoption is even more important for harder-to-scale programs involving the distribution of cash or in-kind aid.
- **Commercial adoption.** An organization explores a potentially profitable product or service, which commercial organizations can then adopt and expand.
 - The **Good Food Institute** works to expand the market for plant-based and clean meat, supporting companies and innovation by connecting experts to opportunities. A project supported by GFI, Counterfactual Ventures, aims at creating for-profit start-ups in the field of clean meat.
- **Mission achievement.** Once an organization has reached a clearly defined, achievable goal, it then winds down its activities. The organization may also pivot if there's another problem it can effectively tackle with its resources and knowledge.
 - Recognizing that fundraising work was no longer neglected within the EA community, the team behind Charity Science Outreach **wound down the project** and shifted their focus toward Charity Science Health and ultimately, Charity Entrepreneurship.
- **Sustained service.** Although this tends to be the default, sustained service is only appropriate if the public or private sectors cannot meet a need. In this case, a nonprofit organization fills the gap, and must constantly build on the efficiency of its program.
 - The Nigerian Government is currently not able to fund widespread cash transfers for vaccinations and the private sector cannot operate a sustainable business model in this field. The charity **New Incentives** will continue to serve as many beneficiaries as possible.

Now that we have a sense of what an endgame might look like, let's look at how we can incorporate the concept into our work. For nonprofits, Gugelev and Stern outline three basic imperatives:

- **Define your endgame early.** Having a clear path forward will keep your organization on track for impact. Working on your endgame will also help refine your theory of change.
- **Focus on your core goals.** Ensure that your organization's activities move you toward your endgame.
- **Prepare your team.** As a nonprofit, your responsibilities are first and foremost to your beneficiaries. But you're also responsible for your employees. Unless your organization's endgame is sustained service, its budget should level off or shrink – this has implications for your staff.

All too far off for you? Many funders (especially in  **venture philanthropy**) will only support you if you have a clear path toward an endgame. “We’ll just grow indefinitely as an NGO” is usually not an answer they like to hear. So **thinking about the endgame has a direct impact on your startgame.**

Gugelev and Stern's exploration of the endgame provides a valuable lens through which to view our work. And as well as allowing us to clarify our impact, it acts as an important reminder. **The ultimate goal of nonprofit work is our own obsolescence.** We dream that the disease we're fighting will be eradicated; that no animals will be born into brief and pain-filled lives on factory farms. Such goals are our lodestar. Reflecting on our endgame brings them back in focus. So yes, the problems that we face are huge, but we can't lose hope because of it – and here, centering our endgame helps. In asking, *How do we scale up?* we spotlight the problem. *What's our endgame?* spotlights its solution.

VII. Problem-solving

Every great project and decision will require problem-solving skills. A large and consistent part of entrepreneurship is getting good at solving a wide range of problems as they come up. Like duct tape, your problem-solving skills will get you out of a jam and temporarily fix many issues as you create space for more thorough and long-lasting solutions. Aspects of problem-solving can be used for almost every decision or challenge encountered. In this section, we will cover:

- Five key steps in problem-solving
- How to come to better solutions by playing steelman solitaire

1. Problem-solving in Five Steps

Author:

Joey Savoie,
Co-founder & Director of Strategy

When you found a charity, problems are inevitably going to come up. It could be that a donor you were counting on does not fund the project, a key employee leaves the job, or a government agency requires a document you have never heard of.

Underneath the shiny website and carefully branded social media pages, most organizations have a consistent stream of diverse and novel problems. A big part of being a well-rounded and talented charity founder is the ability to solve these novel problems effectively. Hundreds of skills and heuristics can lead to better problem-solving: this article covers five of them.

What to do when a problem is detected:

1. Determine how important the problem is
2. Generate solutions

3. Compare solutions
4. Implement solutions
5. Upstream problem-solving

Determine how important the problem is

Whenever a problem is detected, the first step is determining how important it is. There will always be more problems than time to solve them. As such, you have to prioritize how much time and even which problems to solve. Maybe your computer mic is broken and you have to find headphones every time you Skype. Irritating, but hardly an important problem if your organization just got information it might not be having an impact.

The more important the problem, the sooner you should get round to it and the more time you should dedicate to solving it. Often it's worth putting more time into problem-solving than your intuition suggests, as problems get bigger when ignored. (Check out our [productivity article](#) for an additional framework on identifying high-impact tasks.)

Generate solutions

Every problem has a solution. In fact, every problem has multiple solutions. For each problem, you should try to generate a number of solutions. You don't want to give up on a problem without coming up with and attempting multiple solutions.

Often solutions will come to mind quickly, but sometimes it might take some out-of-the-box thinking. Say you're trying to figure out what to do with an underperforming employee. Maybe a couple of ideas come up right away: a) You could fire the employee. b) You could ignore the problem. c) You could do a performance review on the employee and help them work through the issues. These might seem like decent solutions but dig a little deeper and far more come up.

One way to generate more options is to look at the problem from different perspectives. For example, you could frame the problem in a different way. "Why is this employee underperforming" might lead to some solutions you have already thought of. Well, maybe they hate the whole workplace and are going to move on soon in any case, thus voiding the problem. But it can also generate other possibilities. Maybe the employee is underperforming because they do not like their

manager, so maybe the problem could be solved by moving them under a different manager – this is a new solution that could be helpful.

Another way to look at the problem from different perspectives is to ask others, particularly if you combine strategies and use different framings of the problem when asking them. Everyone comes at problems a little bit differently, so one of the big advantages of having a strong advisory board is to come to them with problems and see if they have different solutions. Sometimes it will come from thinking differently: “Is the employee missing a key skill you assumed they had? Maybe they are underperforming due to a lack of training.” Other times it might come from just having different information: “Maybe you should tell the employee to apply for X new job I know about – it seems like a good fit for their skill sets.”

After you have a solid list of possible solutions, you can move to comparing solutions.

Comparing solutions

How much time to spend comparing solutions depends on the magnitude of problem-solving. But keep in mind that generating and comparing ideas is likely worth 10% of the time the solution will take to implement.

When comparing solutions, make a **comparison spreadsheet** or a simple pros and cons list. Remember that solutions are not exclusive – your plan might involve trying one low-cost solution first then implementing another if that doesn't work. By the end of your comparison, you should feel comfortable talking to an advisor or trusted friend (and often this is a good thing to do) about the options you considered and why you came to the solution you chose.

Implementing solutions

The first step in implementing a solution is communicating the plan clearly to everyone it involves. In some cases this might be no one, but in most cases you will have to explain to your team why you have gone with this plan (good thing you have practice from when you explained your solution to your friend). The solution is best presented clearly, assertively and confidently. You do not need to always be confident you have made the right call – everyone makes mistakes. But if you are

confident in your process you can confidently say “I think this is the best call on it given the time I was able to put into it”.

Many solutions require habit-forming. There are great resources on how to do this, including the book *Atomic Habits*. Many require changing your team’s mind on key issues (the best book on this is *Switch*). This can take time and patience. Most of the time a solution will not be universally liked by everyone, but over time and with encouragement, people will adapt.

For example, say you are getting your team to switch to a new task management system. You might plan to speak about it at a meeting your team has regularly. Start with the “need”, or why this is happening. “As you all know, our team has been struggling to stay on top of all the work we need to do. A few important things have fallen through the cracks.”

Next, you want to briefly explain your process in coming to the solution, drawing the person along on the journey you went through. “I talked to a few other organizations and spent some time researching possible solutions. One thing that came up was the type and way our task management system is used. I compared about a dozen task management systems and it seems as though quite a few of them would give our team big benefits.”

Finally, explain the transition and implementation of the idea. Often small steps are good. “We are going to try this new system for a week with our communications team and reevaluate after that. If it seems to help some of our problems, we will switch everyone over.” To help ease the change, document the new processes by writing standard operating procedures: this way everyone is on the same page about how things should be done.

Acknowledge possible flaws or complaints but be reassuring and firm and remind them of the benefit (particularly the long term benefits). “I know it will be a pain to move all our old tasks over, but once we transition we will get more done every week. And the more we can do with our limited resources, the better we can achieve our goals as an organization and save lives.”

At this point, you might think your problem is solved and you can move on. But you can’t quite yet.

Upstream problem-solving

The last step is solving the problem upstream. Solving problems can be categorized in two ways: upstream problem-solving and firefighting. So far the process we have talked about is mostly firefighting steps. Upstream problem-solving refers to actions that can be done to prevent problems from arising, like getting a flu vaccination to lower your chances of getting the flu. Firefighting, on the other hand, refers to reactive actions taken after the problem occurs, like going to a doctor after you feel ill. Most often an ounce of prevention is worth a pound of cure, and the more problems you can solve upstream the easier it will be to fight the (still inevitable) problems that come up downstream. Often you have to experience the flu to know the importance of preventing it the next time. For every problem you solve, you should think: “Does this have a chance of coming up again? How can I prevent this sort of problem from reoccurring?”

A lot of upstream problem-solving comes from the realization of recurrent patterns and the seeing of problems as “another one of those” instead of a unique situation. Of course, each situation is somewhat unique but that does not mean the same underlying problem can be causing it.

For example, say you are looking at three problems on your desk. One is that your operations team is falling behind on their monthly bookkeeping. The second is a promising new idea that you would love to pursue, but you don’t have any staff to spare for it. The third is that your communications team is not responding well to emails sent to your website; many of them seem to fall through the cracks.

These may seem like unrelated problems in unconnected departments, but they could also be a sign you have not built in enough **organizational slack** to deal with a changing and dynamic workload. A firefighting solution might be to pull an all-nighter responding to the emails and doing the bookkeeping. An upstream solution might be to hire more staff or narrow the scope of your project, freeing up organizational time.

Upstream problem-solving does not come naturally. It’s not trained in school and not practiced in most workplaces. There are few resources aimed at teaching upstream problem-solving, but **Upstream** by Dan Heath is a deep guide both to the barriers – i.e. the mindset that a problem is unavoidable, “not my problem”, or can’t be dealt with right now – and to some solutions. Key questions to ask in upstream problem-solving are:

- How do we bring together the right people?
- How can we change flawed systems?
- What's the leverage point?
- How can we spot problems before they happen, steering clear of false positives (e.g. a screening test incorrectly shows that someone has cancer)?
- What does success look like and how do we avoid “ghost victories” (i.e. failures in the trappings of success)?
- How do we avoid causing harm in flow-through effects?
- As nothing happens, who foots the bill?

More materials including a [summary](#) and the [first chapter of the book](#) are available on the author's [website](#).

Summary

When a problem is detected, you want to:

1. Determine how important the problem is: Either dedicate time and energy to it or deprioritize it accordingly.
2. Generate solutions: Use tricks like reframing the question or ask people with different ways of viewing things, or for best results use both.
3. Compare solutions: Use a simple pros and cons list or a more complex spreadsheet to evaluate the options you come up with.
4. Implement solutions: Communicate clearly and firmly the process and solution you came to and set up organizational habits to ensure the solution happens long term.
5. Upstream problem-solving: Consider the cause or origin of the problem and how your organization could prevent future problems of the same type.

2. How to Come to Better Decisions by Playing Steelman Solitaire

Author:

Kat Woods,
CE Co-founder

I have a tool for thinking that I call “steelman solitaire”. I have found that it comes to much better conclusions than doing “free-style” thinking, so I thought I should share it with more people. In summary, it consists of arguing with yourself in the program Workflowy, alternating between writing a steelman of an argument, a steelman of a counter-argument, a steelman of a counter-counter-argument, etc. (I will explain steelmanning later in the post; in brief, it is the opposite of a strawman argument, in that steelmanning presents the strongest possible version of an opposing view.) In this blog post I’ll first explain the broad steps, then list the benefits, and finally, go into more depth on how to do it.

Benefits

1. **Structure forces you to do the thing you know you should do anyway.** Most people reading this already know that it’s important to consider the best arguments on all sides instead of just considering the weakest on the other. Many already know **that you can’t just consider a counter-argument then consider yourself done**. However, it’s easy to forget to do so. The structure of this method makes you much more likely to follow through with your existing rational aspirations.
2. **Clarifies thinking.** I’m sure everybody has experienced a discussion that’s gone all over the place, and by the end you’re more confused than when you started. Some points get lost and forgotten while others dominate. This approach helps to organize and clarify your thinking, revealing holes and strengths in different lines of thought.
3. **More likely to change your mind.** As much as we aspire not to, most people, even the most competent rationalists, will often become entrenched in a position due to the nature of conversations. In steelman solitaire, there’s no

other person to lose face to or to hurt your feelings. This often makes it more likely to change than your mind than a lot of other methods.

4. **Makes you think much more deeply than usual.** A common feature of people I would describe as “deep thinkers” is that they’ve often already thought of my counter-argument, and the counter-counter-counter-etc-argument. This method will make you really dig deeply into an issue.
5. **Dealing with steelmen that are compelling to you.** A problem with a lot of debates is that what is convincing to the other person isn’t convincing to you, even though there are actually good arguments out there. This method allows you to think of those reasons instead of getting caught up with what another person thinks should convince you.
6. **You can look back at why you came to the belief you have.** Like most intellectually-oriented people, I have a lot of opinions. Sometimes so many that I forget why I came to hold them in the first place (but I vaguely remember that it was a good reason, I’m sure). Writing things down can help you refer back to them later and reevaluate.
7. **Better at coming to the truth than most methods.** For the above reasons, I think that this method makes you more likely to come to accurate beliefs.

The broad idea

Strawmanning means presenting the opposing view in the least charitable light – often so uncharitably that it does not resemble the view that the other side actually holds. The term of steelmanning was invented as a counter to this; it means taking the opposing view and trying to present it in its strongest form. This has sometimes been criticized because often the alternative belief proposed by a steelman also isn’t what the other people actually believe. For example, there’s a steelman argument that states that the reason organic food is good is because monopolies are generally bad and Monsanto having a monopoly on food could lead to disastrous consequences. This might indeed be a belief held by some people who are pro-organic, but a huge percentage of people are just falling prey to the naturalistic fallacy.

While steelmanning may not be perfect for understanding people’s true reasons for believing propositions, it is very good for coming to more accurate beliefs yourself. If the reason you believe you don’t have to care about buying organic is because you believe that people only buy organic because of the naturalistic fallacy, you might be missing out on the fact that there’s a good reason for you to buy organic because you think monopolies on food are dangerous.

However – and this is where steelmanning back and forth comes in – what if buying organic doesn't necessarily lead to breaking the monopoly? Maybe upon further investigation, Monsanto doesn't have a monopoly. Or maybe multiple organizations have copyrighted different gene edits, so there's no true monopoly.

The idea behind steelman solitaire is to not stop at steelmanning the opposing view. It's to steelman the counter-counter-argument as well. As has **been said by more eloquent people than myself**, you can't consider an argument and counter-argument and consider yourself a virtuous rationalist. There are very long chains of counter^x arguments, and you want to consider the steelman of each of them. Don't pick any side in advance. Just commit to trying to find the true answer.

This is all well and good in principle but can be challenging to keep organized. This is where **Workflowy** comes in. Workflowy allows you to have counter-arguments nested under arguments, counter-counter-arguments nested under counter-arguments, and so forth. That way you can zoom in and out and focus on one particular line of reasoning, realize you've gone so deep you've lost the forest for the trees, zoom out, and realize what triggered the consideration in the first place. It also allows you to quickly look at the main arguments for and against. Here's a **worked example** for a question.

Tips and tricks

That's the broad-strokes explanation of the method. Below, I'll list a few pointers that I follow, though please do experiment and tweak. This is by no means a final product.

- **Name your arguments.** Instead of just saying "we should buy organic because Monsanto is forming a monopoly and monopolies can lead to abuses of power", call it "monopoly argument" in bold at the front of the bullet point then write the full argument in normal font. Naming arguments condenses the argument and gives you more cognitive workspace to play around with. It also allows you to see your arguments from a bird's eye view.
- **Insult yourself sometimes.** I usually (always) make fun of myself or my arguments while using this technique, just because it's funny. Making your deep thinking **more enjoyable makes you more likely to do it** instead of putting it off forever, much like **including a jelly bean in your vitamin regimen** to incentivize you to take that giant gross pill you know you should take.

- Mark arguments as resolved as they become resolved. If you dive deep into an argument and come to the conclusion that it's not compelling, then mark it clearly as done. I write "rsv" at the beginning of the entry to remind me, but you can use anything that will remind you that you're no longer concerned with that argument. Follow up with a little note at the beginning of the thread giving either a short explanation detailing why it's ruled out, or, ideally, just the named argument that beat it.
- Prioritize ruling out arguments. This is a good general approach to life and one we use in [our research at Charity Entrepreneurship](#). Try to find out as soon as possible whether something isn't going to work. Take a moment when you're thinking of arguments to think of the angles that are most likely to destroy something quickly, then prioritize investigating those. That will allow you to get through more arguments faster, and thus, come to more correct conclusions over your lifetime.
- Start with the trigger. Start with a section where you describe what triggered the thought. This can often help you get to the true question you're trying to answer. A huge trick to coming to correct conclusions is asking the right questions in the first place.
- Use in spreadsheet decision-making. If you're using the [spreadsheet decision-making system](#), then you can play steelman solitaire to help you fill in the cells comparing different options.
- Use for decisions and problem-solving generally. This method can be used for claims about how the universe is, but it can also be applied to decision-making and problem-solving generally. Just start with a problem statement or decision you're contemplating, make a list of possible solutions, then play steelman solitaire on those options.

Conclusion

In summary, steelman solitaire means steelmanning arguments back and forth repeatedly. It helps with:

- Coming to more correct beliefs
- Getting out of unproductive conversations
- Making sure you do epistemically virtuous things that you already know you should do

The method to follow is to make a claim, make a steelman against that claim, then a steelman against that claim, and on and on until you can't anymore or are convinced one way or the other.

VIII. Measurement & evaluation (the tape measure)

“Measure twice and cut once” is a common phrase used when working on a carpentry project. It is equally applicable to running a charity. Without careful measurement and thoughtful evaluation, a charitable project can be off by more than a few inches. It can fail to have an impact altogether.

Measuring your impact and progress both before and during your project is crucial to success. Yet it is a step that the average decision-maker commonly skips. In this section, we will cover:

- Incomplete metrics and how to fix them
- Impact evaluation
- On-going monitoring and evaluation
- Experimental and quasi-experimental design
- Implementing surveys in developing countries

1. Six Examples of Measuring Incomplete Metrics... And How to Fix Them

Author:

Peter Hurford

If you want to create an effective charity, how do you know if it is effective?

A very common answer is, by doing detailed measurement and evaluation. But this is harder than it sounds. While research is hard to get right in general, the difficulty starts just with the measurement. Which metric should you measure?

The Case of the Hard Worker

Imagine that you work for an EA organization. You work really hard. In fact, last week you completed 154 **pomodoros** of work. Was it effective?

It's hard to know without knowing what the pomodoros were actually spent on. Sure, your inputs matter some and it's possible that more pomodoros means more good happening. But instead you need to measure your outputs. Time spent is just a means to accomplish certain high-impact goals. How much good did you accomplish with those pomodoros?

The Case of the Growing Organization

Imagine that you're the CEO of an EA organization. Your EA organization was founded in 2011 and back then you had a budget of \$30K. Now it's 2016 and you have a budget of \$300K. That's 10x growth in five years! That's amazing! ...but is it effective?

The amount of budget your organization spends, like the amount of time your employees work, is only a means to an end. Again, what matters is the results of your budget and what you accomplish with them. It's possible to do more good with \$30K than \$300K.

The Case of the Fundraiser

Imagine that I'm a fundraiser for GiveWell top charities. How do I know whether I'm effective? What if I told you that I was able to send out over one million cold contact fundraising emails in 2015? Would you think I'm an effective fundraiser?

Sure, it's likely that emails lead to donations and that more emails likely means more donations. But without tracking any sort of metric of whether people are (a) opening the email, (b) clicking the donate link, and (c) actually making a donation through that link, it's impossible to tell. Only by measuring actual donations made can we truly measure our impact.

The Case of the Charity Club

Imagine I created a college group called "Charity Club" where we had monthly meetings about donations and career choice. In 2015, we got 100 new members and

held ten meetings with an average attendance of thirty. Does this mean “Charity Club” was effective?

Sure, people attending meetings about donating to effective causes is likely a good thing. But attending meetings about donations is not good in itself. Instead what matters is people actually making effective donations. So we instead should measure the actual donation habits of club members.

But what if club members were just more likely to donate, entirely irrespective of the existence of the club? Ideally we would measure counterfactual impact by randomly only taking some of the interested people into the club and measuring their donation habits relative to the people who weren’t allowed in the club... But chances are that this RCT is not very practical.

The Case of Immigration Reform

Imagine that I think improving immigration in the US is important for economic growth and the welfare of immigrants. So I set up an advocacy website that encourages people to write to their congressperson and encourage immigration reform. How do I know if this website is effective?

One way is that I can measure web traffic. More traffic should be good, right? But what if I get a lot of visitors but no one follows through and writes to their congressperson?

Okay, that’s bad, so maybe I should measure the number of letters to congress people that get delivered. That does measure our influence over the public process, but what if the petitions get ignored? How do we know our petitions lead to legislation change? What if the legislation would have changed anyway?

What we really want to measure is counterfactual legislation change. To do this, we construct an RCT where we randomly select some legislators to be targeted and some not to be and then we see whether the targeted legislators are more likely to sponsor immigration reform than the non-targeted legislators.

While web traffic – or even the amount of letters sent – is a positive thing and might contribute to more immigration reform, it could easily not be connected to immigration reform. Only measuring the right thing helps us check.

The Case of the Developing World Charity

Imagine you're the executive director of a charity that does unconditional cash transfers to the global poor. How do you know you're doing good?

The problem is there is nothing inherently good about people having more money than they used to. Money is just a means to an end. So to know how much good we are doing, we need to see what is happening with the money that we give. What are people spending on? Does the money actually make them happier? To do this, we need to measure the effects of giving money, hopefully with an RCT.

Why might charities focus on incomplete metrics?

In each of the above examples, we consider the metric being measured to be incomplete, or that the metric actually needs to be investigated further before we can clearly connect it with positive impact.

But why might charities focus on these incomplete metrics?

It is easier

Some metrics are way easier to measure than others. It's far easier to measure the web traffic to your advocacy website than to do an RCT on your legislative impact. It's even much easier to measure web traffic than to measure the actual amount of petitions sent. Thus increases in web traffic get cited a lot as a criteria for success, even when it may not be connected to the charity's real goals.

It looks more impressive

Reporting on several metrics looks more impressive because you are showing more data even if the data is not reflective of the good your charity is doing. Additionally, the more metrics you have, the easier it is to cherry-pick the ones that are going well and play down the ones that are not going as well. This kind of practice makes your organization more appealing to donors and members, even if it is ultimately an illusion.

Charities are unsure what the important metrics are

This is true particularly with younger charities, as well as charities with less of a clear focus. When the goals are not clear, this will often result in reporting on several unhelpful metrics or missing very important ones.

Ways to avoid this mistake

Think of the number one most important metric

This is hugely important as it makes clear what your organization is really aiming to do, and how you will measure it. Letting the public know your most important metric also allows them to focus on what really matters.

Be sure of the connection between your metric and to real good happening in the world

Even with a straightforward metric, you have to make sure that it really translates into good getting done. With money raised you would need to look at the charities you are moving money to, and make sure they will accomplish good with extra donation. For website traffic, you would have to make sure website traffic really correlates with actions that you really want to achieve and furthermore make sure those actions correlate with more impact.

Think if it would be possible to cheat this metric

Is it possible to “game” this metric, thus making it less valuable? For example, if I wanted to gain a bunch of website traffic, it would be quite easy for me to invest in non-targeted online ads or just directly buy “views” to my website. Although this would boost my website traffic, it’s very unlikely to cause any real good on the metric I really care about.

Watch out for **counterfactuals**

An easy mistake to make is to measure metrics that have many possible causes. Given that many organizations are working toward the same goals, it is necessary to be able to isolate the impact your organization is having when compared to the wider movement.

Be cautious of longer causality chains

Consider an unconditional cash transfer charity. Their “chain to impact” looks like this:

We give grants of unconditional cash transfers to the global poor → the global poor spend the money on what they desperately need → They are happier because they could afford a basic necessity or invest in their future → Good is achieved.

Furthermore, we're pretty confident about each link on this causal chain because there are multiple studies supporting each link.

Now consider an organization that fundraises for the unconditional cash transfer charity and cites web traffic as their metric of success:

Website traffic → People are then more interested in donating → More people go to the cash transfer charity website to learn more → More money is donated → The transferred cash is spent on basic necessities or investments → People are happier → Good is achieved

Not only is this chain longer, but there is also a huge problem in the assumption that website traffic results in more donations. While we can easily track website visits, it's very difficult to track how many of these visits translates into more donations, and it's easy for the metric to get cheated by getting large amounts of "lower quality" traffic. Generally the more steps you have the more confidence you need to have in each of the steps working.

Focusing on the right metrics in charity entrepreneurship

For folks like us interested in creating the most effective charities, we also need to be careful about metrics. We care most about having the largest counterfactually positive impact on global well-being (for both humans and nonhuman animals). We don't want to look at any incomplete metrics, like the number of people we help, the size of our budget, or how many people read our blog posts. But well-being isn't a very precisely defined metric and very few RCTs look at this.

Instead, there are many more precisely defined metrics that we could measure, such as the impact of our charity on improving the length of life, on reducing the burden of disease and disability, on improving income, improving subjective well-being, etc.

After a lot of research we have opinions on each of these metrics. While each of the metrics is a lot more nuanced than the ones in our examples and none of them are clear examples of incomplete metrics, we do think that some of these metrics are

more complete than others relative to measuring the ultimate goal of global well-being.

2. What Is an Impact Evaluation?

Author:

Ishaan Guptasarma,
Curriculum Developer & Content Writer

Impact evaluation is intended to answer the question: How did outcomes (e.g. the participants' well-being) change due to the intervention?

Thinking counterfactually

One of the trickiest parts of an impact evaluation is thinking **counterfactually**. A “counterfactual” is something which is not fact, but could have been. A counterfactual evaluation is one in which you take the world in which your intervention does not take place, the world in which your intervention does take place, and compare them.

Let's practice applying this thinking to a few different contexts.

Counterfactual thinking: Individual action and market elasticity

Many effective altruists abstain from animal products. How does the world look different, when you abstain from meat, vs. not?

Peter Hurford has tried to answer this question in the article: **How Much Suffering is in the Standard American Diet?**

A good way to start would be to take the equation

Pounds meat consumed / pounds meat produced per animal = number of animals used.

This number is a reasonable estimate of how many animals *you personally* would have eaten or otherwise utilized. But that's not really the number you care about – presumably what you care about is how many animals you've *prevented* from being used.

Peter gets closer to the number we *really* want by adding a multiplier for the

elasticity of demand of each animal product. Basically, this means that whenever you abstain from purchasing meat, there is more meat for everybody else because your non-participation in this market has made the meat cheaper. This doesn't mean you haven't had impact, because this also means meat is less profitable and less meat is produced. If elasticity of demand for chicken is .76, when you abstain from eating four chickens, it influences the market such that three fewer chickens reach the shelves – which hopefully means that fewer chickens are hatched in a factory farm.

By multiplying the equation above by elasticity of demand, your number is now closer to your true, counterfactual impact.

Counterfactual thinking: Careers and replaceability

Gregory Lewis attempts to calculate **how many lives a doctor saves**. After some back-of-the-envelope estimates, he guesses that the average doctor saves 760 QALYs, which drops to 450 QALYs after applying the elasticity considerations described above. In other words, his model thinks that if you convince 10 people to become doctors, you only add 6 doctors to the system, because you edge out 4 people from successfully pursuing medicine – presumably due to some bottleneck such as limited medical school spots, residency spots, etc.

Unlike meat consumption, however, doctors are not roughly interchangeable. Perhaps an exceptional doctor's patients do far better, or perhaps an impact-minded doctor seeks specific positions or moves to locations that increase their direct impact. The question now turns to those 4 people who were edged out. "Who didn't become a doctor *because* I became a doctor and took up their spot? How impactful a doctor would *they* have been?" If those people wouldn't have done any good anyway, you don't have to worry about "elasticity" at all. On the other hand, if those people would've been just as good as you, your counterfactual impact might fall substantially. (Sort of – it's actually more complicated than this because those people wouldn't just disappear, they would apply their talent and labor to another job, possibly one which commands similar value in terms of impact.)

This general phenomenon is known as "**replaceability**", and it's not necessarily just for careers. Just as it's common for several applicants to apply to a single job opening, it's common for more than one charity to apply to a single grant, or for more than one organization to tackle a small problem.

While this is rare in practice because there is so much work to be done, in theory, it might even be possible for interventions to crowd each other out. For example, it could be that a government program, another nonprofit, individual actions, or the market would have eventually accomplished some proportion of the impact had you not deployed your intervention – even if they would have done it less well, or taken longer to do it. If it seems like this might be the case, you should incorporate this consideration in your impact evaluation.

Funding counterfactuals and the impact of charity entrepreneurship

Peter Hurford has tried to find **the expected value of creating a GiveWell top charity**. He estimates that taking a 15% chance to start a charity that can produce twice as much impact per dollar as GiveDirectly and can absorb \$8.3m/year of funding that would have *otherwise* gone to GiveDirectly, is equivalent to \$1.3m/year of donations to GiveDirectly. To make these guesses, Peter must take “funding counterfactuals” into account.

When you take donations toward accomplishing a goal, it is often the case that had the money not gone to you, it would have gone to something else. Therefore, when calculating your counterfactual impact, you should consider where your money would have gone if it had not come to you, estimate how much impact it would have had, and subtract that value from your total impact.

Grantmaking counterfactuals and the impact of earning to give

Some people might read the above considerations and begin to sway away from doing direct work and toward donating money instead, because it’s clear that the person who might have otherwise snagged your high-earning job likely wouldn’t have donated to high-impact charities. At first, you might think that donating washes away the complexities of replaceability and counterfactuals. But while earning to give is a great option that you should always consider (particularly if you think global health is the most important cause area), it is not immune from counterfactual considerations either.

GiveWell writes that **broad market efficiency** in the charity sector may play a role in increasing the difficulty of finding opportunities that are both underfunded and highly cost-effective, since organizations that can clearly show evidence of impact are more likely to be well funded. This means that if you counterfactually did not grant money to someone, they might have been able to fundraise a similar amount

of money somewhere else and still do a similar amount of good. Without people doing direct work and creating new charities, there would be fewer *funding opportunities* that allow money given out by donors to have a high impact. As a result, money is sometimes held in wait for future opportunities, or granted out to less cost-effective charities than it otherwise might have been if there were more people and organizations doing highly cost-effective direct work. For causes other than global health and development, this issue tends to be even more common, as most organizations working in other fields have not yet fully developed the rigorous standards of cost-effectiveness analysis that would allow funding to be allocated and used in a highly efficient manner.

When pursuing high-impact careers, a well-considered choice must take into account particulars regarding what the main bottleneck is in your individual cause area – whether it is the funding landscape and the capacity of various organizations in your cause area to tackle the issue, or something else.

Apply equal rigor and don't get discouraged

It often is the case that the more rigor one puts into evaluating something, the less favorable it looks. Some of the ideas we've covered in this section may feel demotivating – it's a lot of pressure to compare yourself to some hypothetical other person who might have done your job or used your funding just as well or perhaps even better than you might have! Uncovering crucial considerations that cut your expected impact down to a half or a third of what you thought it would be can be psychologically challenging, especially when you have put a lot of work into the project.

A common failure mode is that the more you learn about something and the more closely you scrutinize it, the more you might start to feel that what you are working on may not be as impactful as other opportunities. Additionally, the more you learn, the more you are able to conceive of exciting new opportunities and considerations that were previously beyond your horizon. Effective altruists are sometimes especially prone to this, since they often want to do the *most* impactful thing. While this is generally a good impulse, don't give in to this feeling too easily! Pursuing impact and knowing when to change projects is important, but it is often the case that if you could scrutinize those other projects with the same intensity that you are scrutinizing your own, they also would look less impactful than they might at first glance. Before changing any of your plans drastically due to a negative update, make sure you have **applied equal rigor** in critically examining the alternatives.

You'll always be biased, so get external evaluation

It's natural to be biased regarding your own project. You should try to counter this bias by seeking out external evaluations. Ask some skeptically minded effective altruist who does not have a personal stake in your organization for some unbiased input. Time permitting, you may also be able to get evaluations from us at Charity Entrepreneurship.

Verify your models with real data

In this section, we covered some of the theoretical aspects of an impact evaluation, using examples that allow individuals to estimate the impact of various actions. However, once you start operating on the scale of a medium to large organization, you're likely going to want to evaluate your impact using real data specific to your intervention.

In the [next sections](#), we're going to be covering how to apply counterfactual thinking and equal rigor in the context of the activities of organizations, getting into the details of experimental design and data collection methods.

Summary

An impact evaluation is an attempt to figure out how much better your intervention made the world, relative to if you had not done the intervention. There are many hidden, non-obvious factors that might drastically alter the amount of impact you have. Some of these factors include market elasticity, replaceability, and hidden counterfactuals. There's a tendency for most of these factors to reduce the expected value of an activity once accounted for, so be sure to apply equal rigor to avoid getting unjustifiably optimistic about interventions you know less about and unjustifiably pessimistic about interventions that you happen to know about.

Resources

[Counterfactual considerations](#)

[Replaceability](#)

[Broad market efficiency](#)

How Much Suffering is in the Standard American Diet?

How many lives do doctors save?

Equal Application of Rigor for EA Interventions

3. Monitoring and Evaluation

Author:

Ishaan Guptasarma,
Curriculum Developer & Content Writer

In a previous section, you designed a **theory of change** for your organization. “Monitoring and Evaluation” refers to the process of checking to make sure that the impact outlined in your theory of change is occurring, and that the intermediate steps to that impact are proceeding as planned.

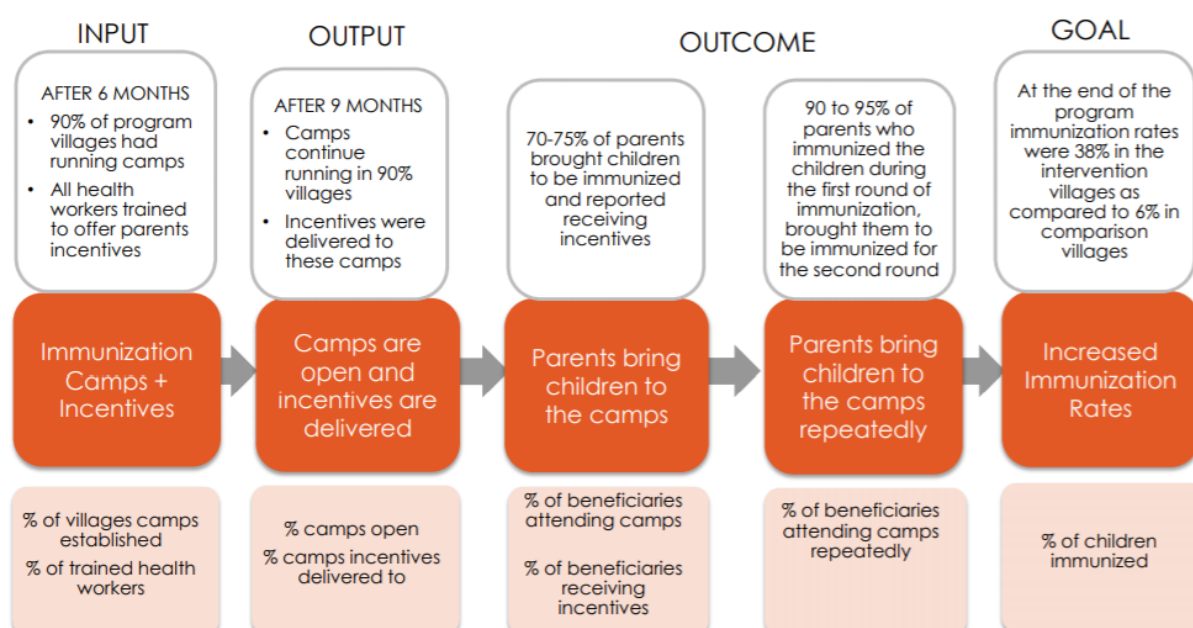
Why do M&E?

The purpose of measurement is to see whether you are having an impact, and to understand the mechanism by which it is occurring so that you can increase your impact. Focus on making measurements that pay out in terms of concrete decisions and changes you might make to your organization’s structure that will help you accomplish more good.

Note that while having an M&E process often looks good to donors and grantmakers, unless you are talking to a very evidence-based organization (e.g. GiveWell) the types of numbers that impress donors are not always the same as the types of data that can help you have a greater impact. Don’t conflate these two – make sure you are clear with yourself on the reasons for a particular M&E activity from the outset.

Building an M&E Process

In your **theory of change**, you specified various kinds of inputs, intermediate steps, outputs, and outcomes. Some of these are measurable.



For each point in your theory of change, think about where you can introduce measurements to give you feedback on whether the changes you have theorized about are happening in practice.

One tool that may be helpful in systematically going through each step in your theory of change is the logical framework matrix, or logframe. Here is a logframe built out of the above theory of change:

Convert to Log Frame

	Objectives Hierarchy	Indicators	Sources of Verification	Assumptions / Threats	
Impact (Goal/ Overall objective)	Increased immunization	Immunization rate (% of children immunized)	Household survey	Adequate vaccine supply, parents do not have second thoughts	Needs assessment ↑ Impact evaluation ↑ Process evaluation ↓
Outcome (Project Objective)	Parents attend the immunization camps repeatedly	% of beneficiaries attending follow up camps	Household survey; immunization card / attendance	Parents have the time to come	
Outputs	Immunization camps are reliably open; Incentives are delivered	% camps open as scheduled; # of kg bags delivered;	Camp administrative data; Random audits	Nurses/assistants will show up to camp and give out incentives properly	
Inputs (Activities)	Camps + incentives are established	% of Camps built, % of camps functional	Random audits of camps	Sufficient materials, funding, manpower	
					Needs assessment

For each step in your theory of change – starting with the direct actions that you take, accounting for intermediate steps done by your stakeholders and third parties, and ending with the final result – you can specify a measurement or source of data to help you confirm that things are going as planned. You can also outline potential factors that might prevent you from impact. Prioritize measures that mitigate these threats and help you make better decisions. Here's a [logframe template](#) that you can use for your own organization. The logframe also comes with a detailed description on how to use it.

Find creative ways to incorporate measurement into your charity

You might imagine that measurement and evaluation mostly consist of rigorous randomized controlled trials. But what if you don't have the time or money for a big experiment? If you randomly select five data points from a population, there is a 93% chance that the population median lies in between the lowest and the highest value. Measurement does not always have to be expensive or time-consuming, and even a small amount of data can be extremely informative. You can find these and other tips in [How to Measure Anything](#), a book about finding creative ways to get figures without spending a ton of money or effort.

Remember that for many purposes, you don't necessarily need highly precise data. You can run a [focus group](#) to get qualitative anecdotal reports from beneficiaries or workers who are more personally in touch with facts on the ground than you might be. You can run surveys that solicit opinions from stakeholders. By choosing your measuring tools carefully, you can get very useful data without spending much time or money.

One arena where you will have a lot of opportunity to measure things is the work conducted by yourself and the senior staff. You might track which projects you spend the most time on and what concrete outputs you produce, or solicit feedback from people who work closely with you. For example, at CE we track the time spent on each stage of our research process so that we understand our internal processes and can quantitatively describe what we are doing to external stakeholders. In addition to helping you quantify your inputs, this can be quite helpful from a [productivity](#) and [management](#) perspective.

Keep data collection decision relevant

It's often tempting to measure everything you can, or to collect data because it seems interesting. However, there's no need to take a measurement if it's not going

to help you increase impact. Unnecessary measurements cost you in terms of design and administration. If you are collecting data from human beings, it also causes fatigue in respondents, making it more difficult to get the data you actually need. For each measurement, make explicit which decision it is helping you take.

Example: Charity Science Health runs a vaccination reminder program. In our surveys, we *don't* need to include a question about the age of the child's parent, because knowing that number will not change any of our decisions. We *do* want to collect data about how many people a surveyor signs up the program per day, because that number can influence *decisions* regarding how many surveyors to hire, show which surveyors to fire or to talk to about underperformance, and help spot scams (for instance, some surveyors may fill in fake data or report an unrealistically high number).

Automation is good, but don't overdo it

Automation is an amazing tool, but be aware that it doesn't always mean less work. Automated processes can sometimes require a long setup process, extended maintenance, and software bugs that will require the attention of skilled and tech-savvy employees. When using automation, try to use preexisting solutions built for exactly the purpose that you are using them and have been shown to work without maintenance, rather than trying to design automated systems from scratch. (Example: use [SurveyCTO](#) rather than trying to build a system from scratch. Most survey tools, including SurveyCTO, build on [Open Data Kit](#) but it is usually worth spending money on more advanced versions with support. For a cheaper SurveyCTO alternative, consider apps such as [doForms](#) but be careful of lock-in and issues down the road as you scale.).

Social desirability bias

Any measurement tool that relies on human reporting is subject to social desirability bias. People will generally follow their own incentives, tend toward answers that put them in a positive light, and lean toward responses that put you and your organization in a positive light in order to avoid insulting you or hurting your feelings. Try to avoid social desirability bias by framing the question so that none of the answers are obviously a "bad" answer. Take answers to questions that are prone to social desirability bias with a grain of salt.

Another way to get more accurate answers is to ensure privacy. Inform people in plain language that no identifying data will be shared with third parties. If you are working in an environment where a door-to-door survey would tend to draw a crowd of onlookers (common in the developing world), give the surveyor explicit instructions to ask the crowd to give some space and ensure the participant can answer sensitive questions anonymously.

Sometimes, you can sidestep this issue by using an objective metric as a proxy for a subjective metric. For example, GiveDirectly has no easy way to figure out people's income without asking them, so they instead use the objective metric of whether the person's house has a dirt floor as one of many metrics. In general, housing characteristics are a common method of estimating poverty in developing world context. This includes materials of the roof, walls, and floor, the cooking method, availability of electricity, plumbing, amongst other criteria.

EXERCISE - SURVEY (10 MIN)

- Come up with ways to ask these questions or measure these metrics in such a way as to reduce social desirability bias
- To farmer: How many hens suffer from untreated broken bones at your farm?
- To politician: Why hasn't your government taxed tobacco at appropriate levels yet?
- To nonprofit manager: How did our training course help your employees become better managers?
- To mother: Have you been taking your child for regular checkups like you're supposed to?
- To potentially illiterate person: Can you read this text message?
- To somebody who might have been influenced by your research: What decisions did you change because of our research?



Assumptions and cultural differences

When interacting with different cultures and peoples who live very different lives from you, you might find that you hold many assumptions about how things work that turn out not to be true.

For example, at Charity Science Health, we had implicitly assumed that people would know their date of birth. When we saw discrepancies in our data, we thought our surveyors were making mistakes. It turned out that people were giving rough estimates regarding dates of birth because they did not know when they were born.

and perhaps didn't necessarily think of that as an important thing to keep track of and be precise about. After that, we pulled dates of birth from official databases of vaccination records. In the same vein, we often found that we were not able to rely on things like maps, street addresses, etc. Even though we were aware that literacy and numeracy was generally low, we didn't automatically realize all the implications that this would have. Sometimes, non-Gregorian calendar ways of expressing dates work better for some beneficiaries (e.g. number of full moons).

Another issue we encountered was working with gender norms which were different from the ones we grew up with. The gender and cultural norms of your respondents and of your surveyors will often influence who feels comfortable speaking to whom, and who feels comfortable going to which locations. We found that female respondents were not always comfortable speaking to an unknown male surveyor, while female surveyors did not always feel safe being sent into unfamiliar locations alone. These types of considerations are highly local, even varying from neighborhood to neighborhood. Talk to your staff and spend some time on site to get a sense of these issues.

Data collection methods

Focus groups

A focus group is a discussion among a small group of typically six to eight participants. Talking things over with a focus group is a great way to get key qualitative background information about how your focus group experiences the intervention. It's often helpful to get this sort of data before investing resources in more rigid and quantitative metrics, as it gives you a basic sense of the relevant and applicable considerations. Without checking first in a focus group, for instance, the metric or question types used in a randomized controlled trial might be biased in ways you would not know about. The general rule here is: talk to many beneficiaries before you run sophisticated studies.

You can use this guide on [How to run focus groups](#) and the accompanying templates in the appendix to get started.

Survey data

Once you get a sense of the relevant questions, quantitative surveys are a great tool to check various parts of your M&E process. Doing surveys in contexts where resources are low comes with a few challenges.

Some common methods of survey data collection include:

- **Phone surveys:** Phone surveys are cheaper and less geographically constrained. It's also easier to monitor surveyors and prevent scams. However, it can be harder to get hold of people because cell phones are often switched off in settings where people can't count on electricity or where data is expensive. Additionally, phones are shared among many people so the wrong person often picks up. Some objective metrics (e.g. does the person live on a dirt floor) are harder to get over the phone.
- **In-person surveys:** In-person surveys are more expensive. It's hard to locate people – many regions do not have signs or street addresses, so you have to ask around to find a specific person. It's also harder to monitor surveyors and prevent scams. However, you can often give longer surveys in person, and get objective metrics.
- **Paper surveys:** Paper surveys typically have low response rates, and require participants to have high literacy. However, on paper, you have the advantage that questions are always asked in the same way, whereas in person or on the phone, surveyors often paraphrase questions in a way that might change the responses. On the other hand, this also means that there is no surveyor to clarify if a participant misunderstands a question. Paper surveys offer more privacy, which may bring more honest answers.
- **Online surveys:** Online surveys are much like paper surveys, but with the added advantage that you can randomize question order, which removes some sources of bias. They're also quicker to pull data out of. Of course, they are limited by respondents having internet access.

Miscellaneous tips regarding surveys in the developing world:

- Be mindful of scams. Build accountability into your system (e.g. if you have surveyors, deploy them in teams and give each team a more trusted supervisor). You can also use tablets to record conversations and see how long was spent on each question. Err on the paranoid side here and implement back checks for as many submissions as possible (e.g. 10%). Frauds you would never expect, such as surveyors filling out fake forms in their hotel rooms instead of in the field, are common. This can be detected

with a combination of GPS tracking, audio recording, resurveying 10% of beneficiaries, and quantitative data checks. Tools such as SurveyCTO are a major asset in this regard.

- Factor into your budget that tablets often break, run out of charge, and get stolen.
- Respondents may not be accustomed to taking surveys and may feel nervous. Put easier and more impersonal questions near the beginning to make respondents more comfortable before delving into potentially more sensitive content.

Some general tips to remember for all surveys

- Role-playing the questions with surveyors beforehand and directly observing the first few responses can help you work out kinks in your process.
- Some people don't take surveys seriously, don't pay attention, respond randomly, or make jokes. To prevent this, include some questions that are designed to weed out unreliable respondents.
- Don't fiddle with analysis methods to make your data look good. Pre-commit to your analysis methods before the survey (sometimes called "preregistration" of your study).
- When taking oral surveys, people remember information given at the start of a list ("primacy effect"), as well as recently spoken information ("recency effect"), better than information in the middle of the list. Because of this, oral survey takers will tend to favor the first and last option that they are presented with.

Running large-scale studies

Many global health, family planning, and mental health interventions will eventually need to run a randomized controlled trial in order to establish evidence that they are working. This will typically be the most expensive part of any monitoring and evaluation project. If appropriate, you should run these once your intervention is running as intended, and you want to find out whether it is worth scaling the project up. You can read more about running and designing large-scale studies in the following article, on [Experimental and quasi-experimental design](#).

Summary

M&E is necessary to make sure that your activities are accomplishing what you set out to accomplish. Running your **theory of change** through the **logframe template** is a good way to augment the process of brainstorming what to measure and how. You would typically start with qualitative **focus groups** to help orient yourself to the realities on the ground. In the intermediate stages, use surveys, internal/external databases (e.g. administrative data such as hospital records), and other forms of data collection to make sure that everything is going as planned. Finally, when everything seems to be working, you might invest in running a **large-scale experiment** to quantify your impact prior to scaling up.

Resources

Related articles

[Experimental and quasi-experimental design](#)

Templates

[Logframe template](#)

External resources

[Goldilocks Toolkit](#)

[Monitoring and evaluation \(M&E\) plan template](#)

[How to run focus groups](#)

[How to Measure Anything](#)

[More Than Good Intentions](#)

[Statistics Done Wrong](#)

[Failing in the Field](#)

[The Goldilocks Challenge](#)

[Logical framework \(logframe\) template](#)

4. Experimental and Quasi-experimental Design

Author:

Ishaan Guptasarma,
Curriculum Developer & Content Writer

When considering an organization rather than an individual action or career, impact evaluation becomes even more important. As the previous sections illustrate, multiple factors can influence how much impact something has, and many of these factors are quite difficult to predict ahead of time.

There are experimental designs that aim to address at least some of these sources of variation. These experiments can sometimes be expensive and time-consuming to conduct, so you should generally begin considering these methods after the intervention is established and you are ready to decide if it should be scaled up. As part of your M&E process (covered in the previous article), you would likely already have implemented easier forms of measurement, such as focus groups, surveys, and internally produced data. The results will have provided reasonable confidence that the intervention could plausibly work, before you consider investing the resources to run a formal study.

Below are just a few examples of experimental designs to give you a rough sense of what is out there – this is by no means an exhaustive list. You can read more about experimental design in [Designing and Commissioning Counterfactual Impact Evaluations](#).

Randomized controlled trials are generally the gold standard for evaluating impact. The study design is straightforward: You randomly assign one part of the population to receive the intervention, and another part not to. These are sometimes referred to as the “treatment arm” (receives intervention) and the “control group” (does not receive intervention). If performed correctly, confounding differences between the two groups disappear, and any observed differences might now be completely attributable to your intervention.

Randomized controlled trials usually generate the highest quality data, because they allow you to clearly understand the counterfactual impact. If there's an economic downturn and your beneficiaries did worse after the intervention, but people who didn't get the intervention did *even worse*, a randomized controlled trial will capture the positive impact of your intervention. If your beneficiaries did better after your intervention, but people who didn't get your intervention also did better and your intervention had nothing to do with it, a randomized controlled trial will capture that too. Without a randomized controlled trial, it can be quite a bit more complex to figure out the true effect of what you did.





Image source: *Impact Evaluation Methods in Public Economics: A Brief Introduction to Randomized Evaluations and Comparison with Other Methods*. A randomized controlled trial cleanly distinguishes impacts from counterfactuals.

Discontinuity regressions take an intervention applied at some arbitrary cutoff, and compare data points on either side of the threshold to determine the effect of the intervention. For example, if one were to tutor every student who scored less than 70%, you could look at the differences in outcomes for those who scored 69% vs. those who scored 71% to get a sense of the effect of the tutoring. This method is a type of “natural experiment” – a situation where two similar groups experience different things, and you quantify the effects of those differences. This method circumvents some of the expenses and ethical dilemmas surrounding random assignment, but is vulnerable to confounding factors. One confounder is that non-linear relationships within the data can be mistaken as a discontinuity attributable to your intervention. For example, if you assume two variables are related linearly (a straight line), but in reality the relationship between them is

described by a logistic function (a curvy S-shape), the inflection point (the middle of the “S”) might be mistaken for a discontinuity.



Left: Image Source: *Regression Discontinuity Design – Andrew Johnston | Economics, University of California*. Discontinuity at the cut off point is used to make hypotheses about the effect of the treatment.

Right: Real-world data doesn't always manifest in straight lines. A logistic curve with an inflection point can falsely mimic a discontinuity, and is an example of how a non-linear relationship between data can create false positives.

Differences in differences is a statistical technique that can be used to compare two groups that differ from each other, but are generally going through the same changes over time. Suppose you have child mortality data from two villages, but are running your intervention in just one of them. The two villages may not be the same – one may have a generally higher child mortality rate than the other. However, the two villages may undergo similar changes over time – for instance, they're likely to be similarly affected by variations in the national economy, changes in the weather, and natural disasters. The name alludes to detecting “differences in differences” over time – if two groups change over time but in similar ways, the method returns a null result. The weakness of this method is that the groups you are comparing may change at different rates for reasons unrelated to the intervention. Remember that a positive result isn't necessarily attributable to the intervention or event that you're interested in, it only means that *something* changed.




Image source: *Designing and Commissioning Counterfactual Impact Evaluations*, a differences in differences design uses the trajectory of a similar group over time to predict what might have happened without the intervention.

Interrupted time series analyses are a method to compare the same population for a long period of time before and after an intervention. For example, because global GDP generally rises at a somewhat predictable rate, you might try using an interrupted time series to check whether a global pandemic had a lasting effect on it a decade later.

This is often a good method to use when there is no possibility of controlled randomization, and no reasonable group that might serve as a natural control. However, there is a risk of unrelated events occurring at the cutoff point and contaminating the data.




Image source: *Interrupted time series regression for the evaluation of public health interventions: a tutorial*. Despite a secular increase in adverse cardiac events, this interrupted time series shows that a smoking ban in Sicily may have had an impact.

Remember, science is messy. It's very easy for factors you did not predict to change everything in your results. It's also very easy to make mistakes with statistics.

Terms to know that might influence or confound your results include

- **“Significance” vs “effect size”**: If a result is statistically significant, it means that what you're seeing might not be due to random coincidence. “Effect size” is how big the effect actually is, and for your purposes, this will generally be what actually matters. (You could have a $p < .05$ result that you reduced malaria in a region by 1% by preventing 500 cases, but that wouldn't be as exciting as a $p < .05$ that you reduced malaria in a region by 10% by preventing 5000 cases).
- **Placebo effect**: You can eliminate some placebo effects using **blinding**. A “double-blind randomized controlled trial” is an experiment where both the participants and the party administering and measuring the intervention are unaware of which treatment arm they are in. If only one party is unaware, it's called a “single-blind”.
- **Correct for multiple comparisons/p-hacking/publication bias**: If you run 100 statistical tests with a 5% chance of a false positive, you'll get on average 5 false positives. Because positive results are often more exciting than negative results, they are more likely to be published. This publication bias means that scientific journals often have a lot of false-positive results. It's also something to watch out for in your own measurements – don't only display those results that come out the way you hoped. You can combat this by preregistering your study before seeing the results, or by statistically correcting for multiple comparisons. Preregistering a study has become standard to avoid looking for correlations until you find one. GiveWell-incubated charity New Incentives, for example, has **preregistered** its randomized controlled trial in Nigeria.
- **Generalizability**: Sometimes, results taken in one context will generalize to another. Other times, they won't. While a randomized controlled trial can demonstrate that an intervention works within a specific context (*internal validity*), you can't be confident that the results will apply in a different region (*external validity*). You will have to try to understand the mechanism by which the intervention works. Read J-PAL's **The Generalizability Puzzle** for more information about this.

More general tips are available from GiveWell's [How we evaluate a study](#).

Organizations to know

If you are implementing a human-focused intervention, you don't have to carry out M&E on your own. There are organizations that can help. Having your work evaluated by a third-party organization well respected in the field can both provide useful data for you and your organization and demonstrate your impact to donors.

In the effective altruism community, [GiveWell](#) often partners with [IDinsight](#). In the broader development community, [J-PAL](#), [IPA](#), and [CEGA](#) have a network of academics who conduct high-quality impact evaluations.

What if a well-designed study says our intervention is low-impact?

Firstly, don't be afraid of shutting down the project if it isn't working. Your time, your staff's time, and your funding are valuable forces for good. When you have clear evidence that there is a way to use resources in a way that can improve the world more effectively, you should stop what you are doing and pursue that instead. Starting an organization is a high-risk, high-reward activity, and it is good to begin each project with plans for both success and failure.

Secondly, don't be afraid of starting a new project. Many of our projects have failed, but we continued to start new ones, because we believed that the *expected* value of our work was positive. If someone offered you a one in six chance to save someone's life for \$100, you would keep taking that bet and rolling that dice over and over again, even if you lost the first couple of times. If you have managed to bring a project all the way to the point of a formal measurement and impact evaluation, it shows that you have what it takes to start over with another project.

Sometimes, you can use the infrastructure you have built to pivot to an entirely new intervention. For example, [New Incentives](#) pivoted three times. They started out using conditional cash transfers to incentivize HIV-positive pregnant women to give birth in clinics, to prevent transmission to newborns. While this was cost-effective, it was difficult to scale. They switched to incentivizing in-clinic births in general, but stopped when there was insufficient proof that the poorly

equipped small clinics would support a reduction in mortality. Finally, they switched again to incentivizing vaccinations. This intervention has been scaled up and is undergoing an RCT, and New Incentives continues to receive [grants from GiveWell](#).

If your analysis shows that you have modest impact, but not enough to be worth redirecting money from the likes of, for example, GiveWell top charities, another potential solution is to seek less counterfactually impactful funding. If your intervention can be funded with money that would not otherwise have gone to something you believe has a higher impact than what you are working on, then the net impact is still positive. Note that in our experience, redirecting funding from outside the Effective Altruism community to organizations considered high-impact according to EA metrics has been challenging (as [Charity Science: Outreach](#), we tried to raise money for GiveWell top charities from non-EA sources), but that doesn't necessarily mean that your organization will not be able to do so.

Summary

Before making the decision to scale up, it's worth investing in some high-quality evidence. Randomized controlled trials provide the highest-quality data, but there are other options if this isn't possible. Some common ones include discontinuity regression, differences in differences, and interrupted time series. [GiveWell](#), [IDinsight](#), [J-PAL](#), [IPA](#), and [CEGA](#) are some organizations that might potentially be recruited to conduct external experiments on your intervention.

Resources

Scientific literacy

[The Generalizability Puzzle](#)

[How we evaluate a study](#)

Study design

[Running Randomized Evaluations: A Practical Guide](#) (book)

[Impact Evaluation Methods in Public Economics: A Brief Introduction to Randomized Evaluations and Comparison with Other Methods](#)

[Designing and Commissioning Counterfactual Impact Evaluations:](#)

(Covers RCTs, differences in differences, and others. Does not cover "Interrupted Time Series" – you can learn about that from the articles below.)

[A Practitioner's Guide To Interrupted Time Series](#)

Interrupted time series regression for the evaluation of public health interventions: a tutorial

Regression based quasi-experimental approach when randomisation is not an option

Organizations

The Abdul Latif Jameel Poverty Action Lab

IDinsight

Other

Data, Economics, and Development Policy 6 courses

Running Randomized Evaluations: A Practical Guide

Mastering Econometrics: Randomized Controlled Trials (online course)

5. Implementing Surveys in Developing Countries

Author:

Erik Hausen,
Senior Research Advisor

This article outlines some simple tips for how to answer questions you have about your program. These tips will hopefully be useful regardless of what stage you are at with your intervention. Most of my experience comes from running surveys in India with J-PAL and Charity Science Health.

Speak directly with recipients

If you are more broadly trying to understand the recipients of your program or how to improve your program, focus groups and directly having more in-depth conversations with those recipients can be much more valuable than a standard survey.

The first place to start when trying to understand your program's recipients should likely be these types of in-depth conversations. You should only move to more standard surveys after having done this.

Start with a phone survey

Phone surveys are cheaper, faster, easier to monitor and should be considered the primary option unless there is a very good reason to do the survey in person. The exceptions to this are when you are asking sensitive questions where you are more likely to get a more representative answer in person (contraceptive use, mental health questions, etc.). Additionally, if you are working in the very poorest communities, it might be hard to identify and reach people by phone.

I recommend initially using Google Forms, Skype and a [Skype Recorder](#). If you start doing thousands of surveys, it might be worth the investment costs to move to [SurveyCTO](#).

Keep it short

It is very easy to spend more time on surveys than necessary. Surveys tend to grow and become more complicated with time. I recommend creating surveys that help you make a specific decision or are focused on a specific topic.

A good rule of thumb is to aim for your surveys to take between 5 and 10 minutes. Following this rule will limit the questions you ask to what really matters and not combine too many topics in one survey. Another way to really get at what matters is to ask, “how does this question inform my decisions?” for every question.

For example, you might be interested in knowing all the demographic information about participants in your program. But unless you actually plan to use that information to change your program, it is likely best to just collect a few basic questions on this.

High-quality data

The quality of the people conducting your surveys should be very high to start with. You should have confidence that the person has the ability to identify issues with the survey independently and can evaluate for themselves if the answers they are getting back from questions are reasonable. If the surveys are being conducted in another language, you will be putting a lot of trust in the person doing the surveys or the person directly monitoring the surveyors. I find the minimum skill required is found in most people coming out of high-level undergraduate programs in Europe or North America. You may want to collect the data yourself if you only plan on collecting a small number of surveys (less than a few hundred).

The risk of making a poor decision because a survey was not designed well or implemented properly is high. The amount of time it will take to set up systems to ensure a survey is working as intended with lower-level staff is significant. If you only plan on collecting fewer than 200 surveys, I recommend that a highly competent member of your staff do the surveys directly with participants. If you're planning to collect more data than this, you may want to hire a lower-level staff member that is directly supervised and monitored fairly heavily by someone with high competence on your team.

Overall, I do not think charity entrepreneurs should be thinking about hiring teams of surveyors until you need to collect hundreds to thousands of surveys every month.

Iterating

You should be very quickly updating on how your survey is doing. The rule of thumb I usually follow for a survey we plan to do fewer than a few hundred times is:

1. Directly monitoring every survey being done and making changes to the survey on the spot until we can do the survey 5 times in a row without me finding a substantial improvement.
2. I will want to look at the data collected after 1–2 days of surveys being completed by a highly competent staff member to see if anything sticks out in the data that indicates the survey is not working as intended.
3. I will then want to quickly look at the data again when we have collected at about the halfway point to check for red flags.
4. Then look at the data after the surveys are completely done and make a decision based on the result.

It is a waste of everyone's time if you look at the data only at the very end and realize that the survey did not collect what you intended to.

Survey software

I recommend using Google Forms and a call recording software for most of your phone surveys. If you need to do very complicated surveys that require more complex logic or you are sending lower-level staff directly to participants' households, I would recommend looking into SurveyCTO.

SurveyCTO is a mobile data collection platform used by large research organizations such as J-PAL and IPA. SurveyCTO will allow you to have complicated survey logic, will track the location of where surveys are happening, and will allow you to audio record the surveys being taken in the field. It also has many other features that ensure your data is being collected as you intended by field surveyors.

I would only look at SurveyCTO if there is something you can't do with Google Forms.

Summary

- Speak directly with recipients.
- Start with a phone survey.
- Keep your surveys as short as possible to start with.
- Have highly competent members of your team or yourself conduct the surveys to start with.
- Quickly update based on early data from your survey.
- Use Google Forms unless you need a more sophisticated option.

IX. Task planning (the clamps)

Often in a workshop, you just do not have enough hands. You might need two to secure a project, one to measure, and one to make a key cut. Most people are only equipped with two hands, and as such we need to use tools to help us. Clamps provide the solution by holding a particular aspect in place so that the carpenter's attention can be directed elsewhere. Good task planning holds your tasks organized and consistently allows you to direct your attention to specific aspects of a project without things falling through the cracks. In this section, we will cover:

- Focusing on the highest-impact tasks
- Delegation and not reinventing the wheel
- Use a task management tool
- Boxing your time
- Review your progress
- Work in deep mode
- Reduce or structure meetings
- Not forgetting the other important stuff

1. Achieving Pareto Productivity: Simple Task Management and Productivity Rules that Go a Long Way

Author:

Patrick Stadler,

Director of Communications & Charity Mentor

The art of task management can seem elusive, with monk-like adherents following complex sets of belief systems to arrive at the holy grail of maximum productivity. While advanced users of task management and productivity techniques might indeed beat the average entrepreneur by far, the **Pareto principle** applies here as well: 20% of effort may give you 80% of the benefits. *Pareto Productivity* (trademark request pending as a medium-priority task in Asana...) presents simple task management guidelines that go a long way. So feel free to cancel your 21-day productivity retreat and return the fancy sleep tracking ring. This will get you covered in much less than one **Pomodoro** slot.

Focus on high-impact tasks only

“Besides the noble art of getting things done, there is the noble art of leaving things undone. The wisdom of life consists in the elimination of non-essentials.” ([The Importance of Living](#)) Before jumping into managing tasks, it is key to select only those with high priority – and leave the others undone. In other words, you apply the 80/20 principle to sorting out the small minority of tasks you should actually work on. For a business start-up, this is straightforward and means understanding the needs of the customers and acquiring more of them, as this [Y Combinator talk](#) on managing time outlines. In the case of a charity, it is slightly more complex as the financial resources and the beneficiaries represent two distinct dimensions. In essence, four core task categories directly contribute to a charity start-up’s success:

- a) Fundraising to obtain the financial resources to operate
- b) Applied research and direct feedback from beneficiaries to design a promising program
- c) Running a pilot program with solid monitoring and evaluation to understand your impact
- d) Operating a financially and legally compliant organization

Avoid any tasks that do not fall under these categories at all or only very indirectly. If you have a basic website, for example, updating or redesigning it is not a direct pathway to attracting more grants. Instead, the number of grant applications sent and warm introductions obtained are much more impactful. Similarly, at some point, [desk research has marginal returns](#) and you better talk to potential beneficiaries in the field and run a pilot with a strong monitoring and evaluation component. Finally, you might operate the most effective charity in the world, but if you get into trouble with tax or other government authorities your future is uncertain.

Feel free to change the list of core task categories to make them more applicable to your context, but resist the temptation to go above five categories or make them too broad. If you need an additional perspective to define success factors, take another look at the [most common ways charity start-ups fail](#).

The [Eisenhower matrix](#) suggests prioritizing tasks according to importance and urgency. We extend this model by including the core task categories from above. This forces you to assign each task to a substantial success factor. This is a first

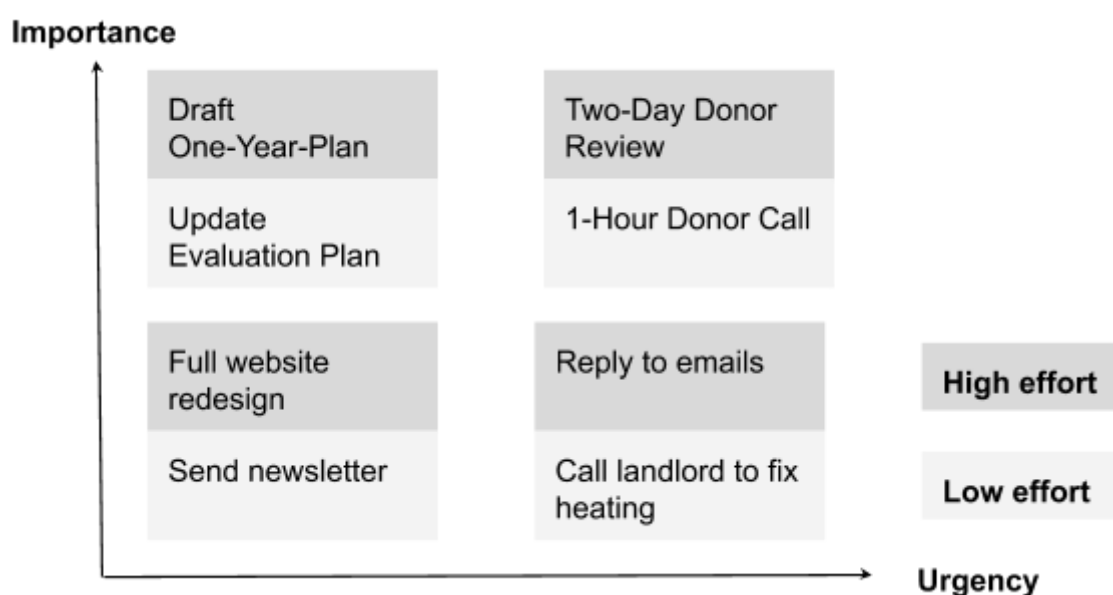
filter against tasks not directly contributing to your charity's success. Moreover, we add effort, as this helps you identify low-hanging fruit.

Core task category	Importance	Urgency	Effort
<ul style="list-style-type: none"> • Fundraising • Research and feedback • Pilot with evaluation • Compliance 	<ul style="list-style-type: none"> • High • Medium • Low 	<ul style="list-style-type: none"> • High • Medium • Low 	<ul style="list-style-type: none"> • High • Medium • Low

Follow these rules as you implement the grading framework:

- Avoid tasks that do not directly contribute to a core task defining the success of your organization.
- Prioritize tasks with high importance and high urgency.
- Do not neglect tasks with high importance but low urgency. Importance trumps urgency.
- Tasks with low importance and low urgency can often be postponed. Tasks with low importance but high urgency are suitable to be delegated.
- Effort is generally less important than importance and urgency. Yet among the important/urgent tasks, you want to prioritize those with the lowest effort first. Pick these low-hanging fruits.

Here is an illustration of sample tasks assessed by importance, urgency, and effort. This assumes that you already confirmed that each task aligns with at least one core task category such as fundraising. You first prioritize importance/urgency. In each cluster, you then give priority to low effort tasks.



Check whether someone else can do it and don't reinvent the wheel

Are you really the best person to implement this task? Your co-founder might be better suited and once you have employees you should try to delegate as much as possible anyhow (with clear task descriptions, responsibilities and deadlines).

Outsourcing is another option that often gets forgotten. No need to spend the weekend reviewing hundreds of field expenses when you can delegate this to a contractor. You can easily find affordable remote freelancers on a platform such as Upwork.com. This works well for tasks such as simple review activities, basic bookkeeping, web research, or IT-related tasks (from developing Google Scripts to updating WordPress). In terms of more expensive contractors such as lawyers, you might be able to find pro bono options (e.g. through [TrustLaw](#)).

If you end up being the one implementing the task, make sure to check for existing advice and templates on the internet. For onboarding, you might consider looking at templates before drafting an [Employee Handbook](#), for instance. Entrepreneurs love to set up things from scratch but often building on existing templates and guidelines can be more productive.

Use a task management tool

Don't be that person who jots down tasks on a random printed out paper – or worse, tries to remember the task without documenting it somewhere. There are

simply too many tasks in the life of an entrepreneur to remember them and it is not the best use of your brainpower.

Using a shared Google Doc or Spreadsheet can be a decent way to track tasks and discuss them with your colleagues. However, this system faces severe limitations too due to the lack of reminders and workflows.

It is best to use a proper task management tool and implement a simplified version of **Getting Things Done** (GTD). As Katriel Friedman of Charity Science Health has pointed out in a talk for CE (unpublished), the key principle of GTD is to avoid “open loops”. These are tasks that are uncategorized, not written down, or without a clear path to completion and therefore may overwhelm and distract you. Create “buckets” that collect all your tasks in a few places (e.g. a notebook and the inbox of your task management app). Place new potential tasks in those buckets immediately, rather than trying to keep track of them using your memory. This way, rather than constantly carrying the mental load of many small tasks, you can review these buckets on a daily or weekly basis. **Daniel Kestenholz** and **Peter Hurford** have written up great summaries of how they use a simplified GTD system in practice.

Here we outline an even simpler form that works for those using a task management app.

How to deal with tasks (email and elsewhere)

Type	Action
Meeting	Add to Google Calendar
Unimportant email	Archive in Gmail
Important email/article	Document it (e.g assign to project folder/label in Gmail or save in related GDrive folder) before archiving
2-Min Task	Implement it immediately
Task (no time to assign)	Keep the task in the inbox of your task management tool, i.e. you write down the task without indicating an assignee or deadline
Task (time to assign)	Assign in task management tool to yourself or someone else and include a deadline. Provide additional context and links if necessary.

As you can see, this list already takes into consideration that many tasks will arrive in the form of emails. Instead of using your email inbox as your to-do-list, you are much more productive if you adopt **Inbox Zero** and move any tasks immediately into your task management app. The Inbox Zero approach also rightly states that you only need to check your email client a few times per day to avoid distraction while implementing your tasks (see Deep Work below).

In terms of tools, you get a discount for **Asana** as an effective altruist organization (through the **EA Hub**). Asana works great for a large organization. For personal usage or side projects, **Todoist** is a strong option. Check out this **review** if you would like to consider different apps. In the end, it is less important which tool you pick but rather that you and your team stick to it.

Box your time

Defining your high-value activities and turning them into tasks is important but not sufficient. Let's say you have to finish a grant application in the next three days. You've got a critical task at hand with a clear deadline. The implementation, however, very much depends on the time needed to complete the task. Timeboxing (also known as **timecapping**) allows you to estimate the required time and book it in your calendar. In the example, you might reserve a slot in Google Calendar one afternoon from 2 pm to 6 pm to finish the grant application.

Timeboxing works for any task: from **research and decision-making** to daily operations. It has a range of advantages. Timeboxing...

- Assigns a concrete time value to your task.
- Helps you avoid overspending time. You realize from the beginning that your time is limited, so you focus and stay within the timeframe. Hence, timeboxing is an excellent tool to force you to implement the 80/20 principle at the individual task level as well. Instead of going off on tangents, you remain focused on the core deliverable.
- Prevents paralysis and indecision. You have clearly defined how long you spend on something. Sure, you might have to update your estimate but this is very different from conducting tasks without any time estimates or deadlines.
- Gives you control. You are the one to define how much time to allocate. This enhances your feeling of **autonomy**, one of the key drivers of job happiness.
- Enhances transparency in your team. Your co-founders see what you are working on and understand why you are busy.

Here are a few best practices in implementing timeboxing:


1. Use one calendar for all your assignments and meetings and share the calendar with your team.
2. Work at least two hours per day on your top goal, ideally more. Pick times when you are most productive, say, in the morning.
3. If you use [Calendly](#) for scheduling meetings, restrict slots to times when you have less energy, for example, late afternoons.
4. Combine similar shorter tasks into one block (e.g. review applications for an intern and a full-time position).
5. Include breaks and logistical slots such as lunch and transit. You are not a robot and your calendar should reflect that.

Review your progress

Each day – and more extensively each week – check your progress on task management by going through this checklist.

1. How much time did I spend on the different tasks?
 - a. For this purpose, you compare the time “boxed” with an estimate of the actual time spent based on your calendar. You planned, for example, to spend three hours on updating your monitoring and evaluation strategy but exceeded that slot by two hours. If you systematically under- or overestimate time for certain tasks that gives you helpful guidance. For example, preparing a 1:1 meeting with an employee usually takes me around 30 not 15 minutes. Advanced users can use a time tracking app such as [Clockify.me](#) to get accurate numbers of time spent per task/category.
2. Is the actual time spent in line with your focus on high-impact tasks? Check against importance, urgency, and effort again.
 - a. You might reassess a task after realizing how time-consuming it is (effort). Other developments at work might have corrected the task’s importance upwards or downwards.
3. Is there nobody that could help with or finish this task for me?
 - a. Based on your progress, you might reconsider delegating a task or getting help from a contractor/freelancer.
4. How do I need to change the deadlines and assignments based on this in my task management app?

- a. The progress made and your reevaluation of the task define whether updates regarding assignee and deadline are necessary.
- 5. What are the implications for timeboxing for the next day/week?
 - a. Finally, the reevaluated task with an updated deadline gets more or less time allocated in your calendar. In the case of an important task for which you continue to be the lead, you might timebox larger slots to meet a deadline. In another case, you might delete a timebox, as you decided to delegate finishing a task.



Work in deep mode

Working on your fundraising strategy, responding to emails while helping out your new colleague over instant messaging. Does this sound like your typical workday? Then you better consider **Deep Work**, as presented in the classic by Cal Newport. The basic message is one that resonates intuitively and has been proven in studies: multitasking and distraction undermine productivity (and flow experiences that contribute to a fulfilled work life).

Deep Work suggests building your whole day around carrying out important tasks without interruption: “Instead of scheduling the occasional break from distraction so you can focus, you should instead schedule the occasional break from focus to give in to distraction.”

Timeboxing, introduced above, is a key tool to arrive at deep work. You should also stop checking your email and turn off notifications from instant messaging such as Slack as you focus on the activity at hand.

Another key concept is productive reflection whereby you give yourself time to think about a certain problem. This is not your typical work task, as it is less linked to a specific outcome such as writing a report. Yet it is also not leisure, as you contemplate a work challenge from various angles. Productive reflection can take place in relaxed settings, say, on a walk or under the shower. For some, this comes naturally. If, however, you find yourself running from one task to the other and lack time for thinking through problems creatively, make sure to dedicate at least one to two hours per week to productive reflection. As an example, your outreach to fish farmers might not have been as successful as hoped for. In productive reflection, you approach the problem from a high level (why do you need to talk to fish farmers? What are all the theoretical ways to reach fish farmers?) and consider different alternatives (what if we set up a hotline instead of sending email newsletters?). The goal is to consider many options in brainstorming mode and follow **first principles**.

Reduce or structure meetings

Meetings are often not the setting to create the building blocks for your charity start-up. When was the last time you created an M&E strategy or fundraising plan in a meeting? As Paul Graham points out in **Maker's Schedule, Manager's Schedule** frequent meetings can interrupt quiet work on the outputs you need to deliver. Nobody would dispute that meetings are essential for some coordination or even some creative problem-solving tasks. But in many work settings, they still take place too frequently and in an unstructured manner.

Here are some basic guidelines for getting the most out of your meetings:

- Apart from weekly team meetings and **one-on-ones** generally don't schedule regular meetings where there is no obvious need for in-person coordination. Team meetings and one-on-ones have an important social component so it matters less if there are no important topics for discussion.
- At least, prepare an agenda before the meeting and document decision-points after the meeting. These decision points can be turned into tasks in your task management app.

- At best, someone prepares options for decision before the meeting and shares them with all participants. An extreme form of this is practiced at Amazon where employees write **multi-page memos** ahead of meetings.
- In terms of meeting scheduling, a best practice is to schedule for 25 and 50 minutes. This allows you to switch location and refresh after each meeting.

Don't forget the other (more) important stuff

This article covers task management and productivity in a relatively narrow sense. The focus is on the immediate work setting and delivering results. While the tools presented here are impactful, more holistic strategies might even be more important. The good news is, you are already fully aware of them. You might just need to commit to implementing them more (see this summary of **Atomic Habits**).

- Sleep well
- Eat well and plant-based
- Spend time with friends and family
- Exercise or at least move (as a workaholic consider Steve Job's famous walking meetings)
- Outsource chores (e.g. see this **time saving assessment** by Joey Savoie)
- Take weekends off and schedule vacations during which you completely disconnect
- Practice mindfulness/meditation (indeed, it would not be an article about productivity without at least one reference to meditation)

As you implement most or some of the practices introduced in this article, you have every right to add the title Pareto Productivity Pro to your business card and LinkedIn profile. You might not yet be an ordained monk in the order of productivity but you are slowly getting there.

Summary

1. **Focus on high-impact tasks only:** Prioritize tasks that give you 80% of the return with 20% of the input. First, make sure each activity is directly linked to at least one of your charity's core success factors (e.g. fundraising). Second, classify tasks according to their importance and urgency. Third, consider effort as an additional criterion to identify low-hanging fruits.

2. **Check whether someone else can do it and don't reinvent the wheel:** Others might be better suited or have more capacity to carry out tasks. This includes co-founders, employees, or freelancers/contractors. The latter can be hired on platforms like Upwork.com and do well at basic IT and quality review tasks.
3. **Use a task management tool:** With a task management app such as Asana or Todoist you can easily implement a basic form of Getting Things Done where you are on top of all your high-priority activities. This ties in well with Inbox Zero, which avoids abusing your email inbox as a task management hub.
4. **Box your time:** Timeboxing puts your top tasks as reserved slots right into your calendar. This forces you to implement individual tasks with an 80/20 mentality to remain within the given timeframe. Other advantages are increased transparency in the team (everyone sees what you are working on) and a sense of autonomy (you are the one to reserve the slot for each task).
5. **Review your progress:** Each day – and more extensively each week – check your progress on task management by going through a simple checklist. The basic idea is to compare the time you spent on a task with its priority and reassess importance, assignee, and deadline if necessary.
6. **Work in deep mode:** Focus on one task at a time and avoid any distractions, e.g. by turning off instant messaging notifications. Plan for at least 1-2 hours of reflection per week during which you ponder a business challenge in brainstorming mode.
7. **Reduce or structure meetings:** Key work usually does not take place in meetings, so keep them limited and – if necessary – structured with a clear agenda (or even preparatory memos) and a decision-making framework as output (who is doing what by when?).
8. **Don't forget the other (more) important stuff:** The most important productivity hacks are closely tied to personal wellbeing. Yes, you know the drill: sleep well, eat well, exercise, ...

Reading list

Internal resources

- [Six Ways Your Charity Start-up Might Fail – and How to Prevent That](#)
- [The Importance of Time Capping](#)

External resources

- [How Timeboxing Works and Why It Will Make You More Productive](#) (HBR)
- [Minimalist time-management system](#) (Daniel Kestenholz)

- [How to prioritize your time](#) (Y Combinator)
- [Increase your productivity using an effective-altruist approach](#) (Tanja Rüegg)
- [Inbox Zero](#) (Merlin Mann)
- [Timesaving Activities](#) (Joey Savoie)
- [Productivity 101 for Beginners](#) (Peter Hurford)
- [How I Am Productive](#) (Peter Hurford)
- [100 Most Useful Productivity Tips](#) (Filtered)
- [Maker's Schedule, Manager's Schedule](#) (Paul Graham)
- [How to Use Gmail More Productively](#) (Andreas Klinger)
- [The Great CEO Within](#) (Matt Mochary)
- [The Complete Guide to Deep Work](#) (Doist)
- [Atomic Habits: Summary](#) (Nate Eliasson)
- [The Best To-Do List Apps for 2020](#) (PC Magazine)
- [First Principles: The Building Blocks of True Knowledge](#) (FS)

X. Independent experts (the cell phone)

No matter how talented you are, you will always need outside help. Like a carpentry expert would call a friend to get a second view, great decision-makers know how and when to get advice from others. External experts, wise advisories, and thoughtful outsiders can provide tools and perspectives that would otherwise be missed and neglected. As with using a cell phone, it is equally important to understand when to ignore certain notifications. In this section, we will cover:

- The pros and cons of speaking to experts
- Our CE process for getting expert views
- How to know which experts to trust

1. How CE Uses Expert views

Author:

Joey Savoie,

Co-founder & Director of Strategy

Charity Entrepreneurship's research process involves **multiple stages** and four methodologies: **cost-effectiveness analysis** (CEA), **informed consideration** (IC), **weighted factor models** (WFM), and **expert views** (EpV). In this document, we explain why and how CE uses EpV as part of our research process, and discuss its strengths and weaknesses as a methodology. .

Garnering expert views consists of speaking to experts who may have broad, domain-focused, or specific knowledge of the field. It is particularly useful because experts can often give a broad overview of a topic, allowing researchers to gain a comprehensive understanding of an idea. However, it is not the only method we rely on because human judgment often suffers from cognitive biases.

CE uses EpV at three stages of our research at different levels of depth: first five minutes, then two hours, and finally twenty hours.

- First, we conduct an **idea sort** to sift out the top thirty or so ideas from our initial list of hundreds. We assess each idea using all four methodologies (i.e., CEA, EpV, WFM, and IC), spending twenty minutes per idea.
- Next we produce a **prioritization report**, spending two hours researching each idea from the idea sort and ordering them according to their promise. We use a different methodology for each cause area: EpV is used for our **family planning** ideas. At this stage we only apply one methodology to ensure sufficient depth.
- Finally, our eighty-hour assessments of the top interventions involve all four methodologies.

Concretely, the lead researchers use EpV by holding one-hour interviews with broad experts at the initial stage, domain experts at the second stage, and a mix of experts with various levels of knowledge, including specific knowledge, for the last stage of the research.

Who are experts?

Speaking to experts is a common way to gain a lot of information about a topic quickly. Experts can synthesize an informally large amount of knowledge into layman's terms that are much easier to understand than a meta-analysis or other form of formal synthesis. When we speak about experts we are referring to three different groups of individuals:

1. Specialist experts
2. Domain experts
3. Broad experts

Specialist experts will often be highly versed but in a very specific situation or content area. For example, a fish disease specialist would fall under specialist domain experts. They can provide a piece of the picture but often not a broad comparison. If your goal is to start a charity that helps the most fish, they would not be able to compare disease to transportation issues, and often specialist experts would not even offer a guess on it. However, they would be able to provide highly specific information about disease rates in a given species and situation that you have identified as promising.

Domain experts are experts who have a sense of a single area. They might know about many different possible factors that affect a single type of fish but would not be able to compare a fish-based intervention to a chicken-based one. Heads of nonprofits in a given area would be good examples of domain experts.

Broad experts can provide comparisons across different domains, for example, a funder who supports half a dozen different fish organizations might have a strong sense of how disease compares to transportation even if she does not have the same detailed sense of specific diseases as the specialist expert.

We see all of these experts as very helpful but in very different situations.

Example conversation notes (GiveWell and Charity Science Health):

A conversation with Joey Savoie, January 26, 2017

Participants

- Joey Savoie – Co-Founder, Charity Science: Health
- Elie Hassenfeld – Co-Founder and Co-Executive Director, GiveWell
- Sophie Monahan – Research Analyst, GiveWell

Note: These notes were compiled by GiveWell and give an overview of the major points made by Joey Savoie.

Summary

GiveWell spoke with Mr. Savoie of Charity Science: Health (CS:H) as part of GiveWell's Incubation Grants work to support the creation of future top charities. Conversation topics included updates on CS:H's plans for testing and scaling up its text message (SMS) vaccine reminder program, the timeline for seeking additional funding, and cost-effectiveness analysis of the program.

Example of synthesized expert data (Charity Entrepreneurship 2018):

Broad area	Intervention	Ordered Sub-interventions	Weighted average	Average	Range in views
Corporate Outreach	Broad area ranking	NA	2.8	2.7	Small
Research	Broad area ranking	NA	2.6	2.7	Large
Product Creation	Broad area ranking	NA	2.5	2.5	Large
Political Outreach	Broad area ranking	NA	2.2	2.3	Small
Welfare Condition Improvements	Broad area ranking	NA	2.5	2.2	Mid
Wild Animal Suffering (WAS) / Bugs	Broad area ranking	NA	2.3	2.2	Large
Improving Existing NGOs	Broad area ranking	NA	1.8	1.8	Mid
Events	Broad area ranking	NA	1.6	1.7	Small
Veg Outreach	Broad area ranking	NA	1.8	1.6	Small
Other	Broad area ranking	NA	2.0	2.0	Mid

Why is this a helpful methodology?

Experts are in many ways the broadest source of information. They rarely give specific conclusions but a broad overview of a large field, covering a lot of ground and have views that are often easy to explain in terms of conclusion but hard to explain in terms of the factors that went into their formulation.

Reasons they are a helpful resource (in rough order of strength):

- **Utilize a large number, and a variety, of evidence sources**
- **Apply common sense filters**
- **Quickly assess weaknesses**
- **Are insensitive to model errors**
- **Use new sources of information**
- **Are a respected source of information**
- **Offer multisession engagement**
- **Can directly compare possible strategies**
- **Provide field-level convergence**

Utilize a large number, and a variety, of evidence sources: Experts have formed views using a wide range of sources of information. They often form their perspectives based on a number of studies, conversations, personal experiences, and other sources. These diverse sources are combined into a single view. This has a number of advantages including making their conclusions more robust and grounded.

Apply common sense filters: When you are new to a field you do not have a strong common sense filter; however, experts have often seen many projects come and go and have a strong sense of things that will be more successful or impactful. Experts often have a sense in their field of what things might go wrong in a project or lead it to failure. These filters can be helpful and informative in prioritization among areas and making long-term plans.

Quickly assess weaknesses: As a result of being able to talk directly to an expert and lay out specific situations and combinations of ideas, it is easier to identify flaws in reasoning or possible areas in which an idea could fail. This information would often be hard to find from informal research or even deeper, more systematic research.

Are insensitive to single-number model errors: One of the biggest concerns with many-multiplication models such as CEAs is that a single error such as a mistyped number can have a large effect on the model. Experts on the other hand are rarely overly affected by a single model or a single number and tend to be slow to update on shocking conclusions. They more intuitively and directly apply “extraordinary claims require extraordinary evidence” heuristics.

Use new sources of information: Experts often have connections and knowledge about what resources are worth considering for further research or which other experts are worth talking to. This methodology, when applied, thus lends itself to finding more information and getting a clear path of who to talk to next or what resources to read. Experts sometimes have access to studies or other research that is not yet available or easily accessible in the public domain. Well-positioned experts can often have access to information a full year before it is publicly available. They also often have details on what studies are being worked on and will be completed in the near future.

Are a respected source of information: Talking to and getting viewpoints from experts is a respected and even expected tool to use when researching an area. It is also common practice in many fields including charity evaluation.

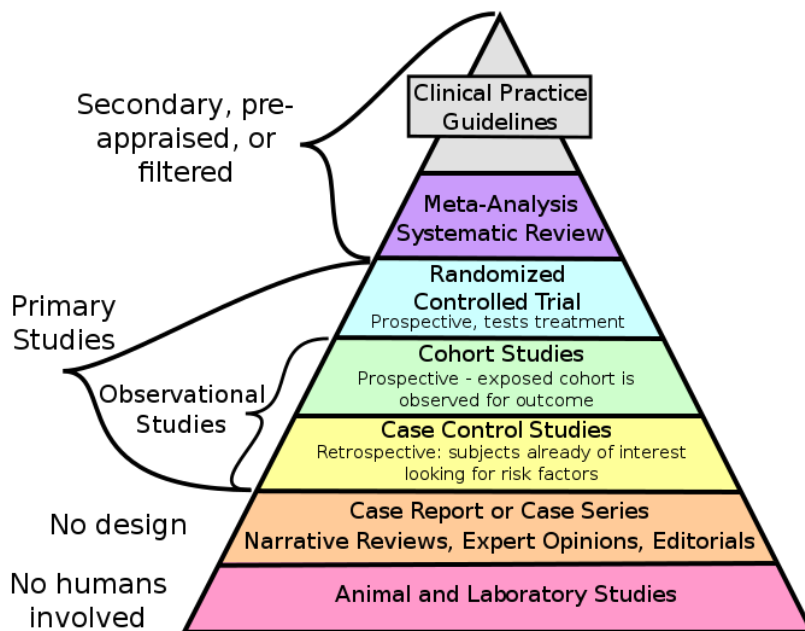
Offer multisession engagement: Experts are one of the few sources of information that can directly engage in back-and-forth discussion, which means that after speaking to them with a more basic idea, you can converse again about the changes or advancements that have happened in your thinking.

Can directly compare possible strategies: Experts can be given highly specific plans and compare different elements of them much more quickly than a more formal model like a CEA can. For example, if you are considering three interventions in three different countries with three different partner organizations, the number of permutations quickly becomes overwhelming for a formal model. Experts, however, can compare multiple iterations and suggest which combination seems to have the highest impact or is the most promising to research further.

Provide field-level convergence: Experts can give a sense of whether many individuals within a field have a fairly unified view on something (e.g., if all three experts you speak to agree on a topic) or if there is a variety of views on a topic (e.g., three experts give three different answers). If an area has a high level of convergence, it is good to get these conclusions, and if it does not, that leaves open more areas that should be considered or researched.

Why experts are not our only endline perspective

Despite experts being a helpful source of information, they are not our only endline perspective. When viewing an evidence table, there is a reason why EpV is often near the bottom.



Experts have a number of weaknesses that have been demonstrated to negatively affect their judgment, and studies have shown in some areas, such as predicting the future, that “many of the experts did worse than random chance, and all of them did worse than **simple algorithms**.” These concerns limit experts’ usefulness and make us confident that they should not be the only perspective used. Many of our biggest concerns with experts are **cognitive biases** that cross-apply to the vast majority of human judgments. Not all the following concerns apply to every expert, but they are generalized concerns that will apply to a large number of experts.

- **Unequal application of rigor**
- **Inconsistent and unclear epistemology**
- **Cognitive biases**, including but not limited to:
 - Anchoring
 - Groupthink
 - Illusion of control (weaknesses with randomness)
- **Lack of transparency in argument generation**
- **Memory concerns**
- **Limited specificity**
- **Lack of decisiveness**

Unequal application of rigor: A major concern with experts is equal application of rigor. Given all the information currently available, a motivated actor can find evidence supporting almost any viewpoint. Thus a fairly weak argument could hold a lot of weight in an expert's view if he or she has not considered it skeptically or if it fits a prior worldview the expert holds. Similarly, if an expert does not like an idea, he or she finds it easy to be significantly more critical of it than would be justified compared to other ideas or viewpoints he or she holds. This rigor concern makes it highly challenging to take expert conclusions without a deeper sense of how they, for example, react to any new idea.

Inconsistent and unclear epistemology: Another factor that makes expert judgment less strong a source of evidence is the relative rarity of a formal or consistent epistemic system. Experts often have views about how to weigh different types of evidence but few have thought about this problem explicitly, and very few have publicly laid out how they would compare and integrate different pieces of evidence into their endline viewpoints.

Cognitive bias: There are a number of cognitive biases that affect humans. Experts are fundamentally just more informed humans and thus generally suffer from the same biases. Some evidence suggests that experts can be affected even more strongly by some biases than the general population. One mitigating factor is that if multiple experts are spoken to, their biases will not necessarily overlap, and their average quality of judgment tends to do better than that of a single expert. There are hundreds of biases that can affect judgment and decision-making, but some that seem particularly relevant to experts when considering charity ideas are:

- **Anchoring:** When an individual depends too heavily on an initial piece of information offered (considered to be the “anchor”) in making decisions (1). Experts can often anchor on a specific idea for a charity early in a conversation or before the conversation has even started. Many experts will have projects they have already supported or invested time into, and these existing projects will generally be compared to any new idea with a high level of comparative skepticism regarding competing or different ideas.
- **Groupthink:** A psychological phenomenon that occurs within a group of people in which the desire for harmony or conformity in the group results in an irrational or dysfunctional decision-making outcome. Group members try to minimize conflict and reach a consensus decision without critical evaluation of alternative viewpoints by actively suppressing dissenting

viewpoints and by isolating themselves from outside influences (2). In the case of charity ideas, if an idea has not been previously tested or considered by experts in the field, often they will be more inclined to dismiss the idea than they would if the same concept was presented by someone connected to their ingroup. Although this is a useful heuristic for experts to use, it can make them underweight new ideas relative to more established ones, particularly if the new ideas are generated using an intelligent process (e.g., CEAs).

- **Illusion of control:** The tendency for people to overestimate their ability to control events; for example, it occurs when someone feels a sense of control over outcomes that he or she demonstrably does not influence (3). This connects closely to experts having difficulty discerning effects from randomness or noise. The way this connects to charity ideas is that often experts will put more weight on personal experiences they have had; for example, if idea A has worked in the past, idea A will always work, and if idea B failed in the past, idea B is likely to fail. These are often assumptions held without careful consideration of the environmental factors that are different or the non results-focused factors; for example, a higher chance of failure might be worth it if the win is several times larger.

Lack of transparency in argument generation: Experts have formed their views using a considerable number of sources and experiences. A byproduct of this is the great difficulty of tracking down the basis for a given viewpoint. This can make it very challenging to confirm or disprove a given idea or even know how much weight it should be given. This is not the fault of the expert but is a flaw inherent in expert-based information.

Memory concerns: Evidence has demonstrated that memory is a fallible tool, but generally when speaking to an expert there is a high level of reliance on the expert's memory. A remembered version of a study or conversation could be significantly different than the original. It is also hard or impossible to detect these memory effects given the lack of transparency.

Limited specificity: Many experts are unwilling to give specific estimates such as a percentage-based chance of success. Experts are often unwilling to make claims that could be used in other methodologies such as CEAs, particularly if those claims cannot be anonymized.

Lack of decisiveness: Similar to the specificity concerns, experts are often unwilling to make decisive claims even when taking a neutral or unsure stance would have its own ramifications. This is often taught practice in academia and can be a good habit when it comes to truth-seeking, although it can impair comparison among different options.

How much weight we give experts

Experts are quite an important source of information, and a lot of fundamental information comes from EpV as interpreted by others. However, experts suffer from many human biases that more algorithmic systems are less affected by. Experts are one of our five perspectives, the others being: WFM, team intuition, CEA, and our prior view. We expect our experts to generally hold more than one-fifth weight but likely not more than one-third of the weight of the model. Given the considerable variation depending on the specific charity idea and cause area, we expect experts to be stronger in areas where it is harder to get solid numbers and commonsense intuitions can be used as an effective guide.

Our process for speaking to experts

Selecting

Who do we count as an expert, and how do we speak to experts? We have two main processes to generate expert names. Our list of experts is generally created opportunistically instead of systematically.

- Finding via research. As we are conducting directed or undirected research, names often come up of key people in the area. These names are then noted down and later contacted as experts.
- By recommendation. Often when we speak to experts we ask for others who would be helpful to talk to. Thus many of our experts come from peer recommendations.

The types of experts we end up speaking to in practice fall into the following categories.

1. **Broad experts:** Generally broad experts are the smallest group we speak to. These are people involved in comparing a range of interventions within a given area. Often their knowledge will be more generalized when compared to domain or specialist experts. An example of a broad expert might be a

researcher at an evaluation organization such as GiveWell or Copenhagen Consensus or an author on the Disease Control Priority reports. However, they could also be a large funder or cross-area researcher who has a good sense of the space overall. These experts would be considered “broad experts” in a cause area such as global poverty.

2. **Domain experts:** We generally speak to many domain experts, and they often consist of the largest group of experts we speak to. Often such an individual could be the author of multiple studies in a given area, someone who worked on a synthesis of the intervention or domain being considered, or someone who conducted an evaluation or meta-analysis of the area or an adjacent area. These experts would be considered “domain experts” within an intervention area such as vaccination reminders.
3. **Specialist experts:** We often speak to a small number of specialist experts. These are individuals who are highly informed but in a small niche within a given charity idea or intervention we are considering. This might be a biology specialist who has deep information about food fortification uptake rates. It could be a person experienced in an element of running an intervention on the ground, such as country-level experience, or an individual who has written a study related to the execution of a given intervention. We consider these experts specialists who could have insight into an element of, but not the whole, intervention.

Contacting

When contacting an expert for the first time, we generally use an email similar to the following.

Dear _____

I am a research associate at an organization that researches and funds new nonprofits to put that research into action ([our website here](#)).

_____ person recommended that I speak to you because of your background in _____ OR I am researching _____ and I read your paper on _____, which drew my attention because [detail]. Based on this, I thought you might know the answers to a few questions about the topic.

Some questions I have are: _____ and _____. Would you happen to have the time to jot down some quick answers to the above, or maybe discussing it via Skype

might be easier? We'd really love to have your input and research inform our funding decisions and what charities to work on in the sphere.

Best regards,

How to speak to experts

Experts are ultimately just people like anyone else, so most standard conversational rules apply to them. A few elements to highlight are:

- Be humble – At the point you are talking to experts, you are new to the field, so you should try to come across as someone who is surprisingly informed for a nonexpert. Ask for only a little bit of time, such as one Skype instead of longer commitments. Say you are happy to keep any comment anonymous if they review the conversation notes, or say it should be so during the conversation. Try to take a broad interest in the topic matter as a whole even if it is not directly tied to the question you are asking.
- Be prepared – Being thoughtful with an expert's time is important. If they have written a whole book on a given topic, for example, you should at the least review a summary before talking to them about it. The same goes for website content they have created. As well as reading content beforehand, you should think about the most important questions and which ones could get cut if you run out of time. Have a backup system for contacting if the first one does not work (e.g., Google phone credit you can use to call someone if Skype is not working).
- Frame opposing views using a citation – If you want to push on a point or perspective that an expert has claimed, do not describe it as "I am skeptical of point A"; instead tie it directly to your research ("A different expert I spoke to was skeptical of point A").
- Go deeper – Try to cover the key questions on your agenda, but if something comes up that seems important you can ask more questions relating to that area. Ask follow-up questions, such as "You said that . . . Why do you think that is?"
- Ask comparative questions – Few experts will have a great sense of what the percentage chance of something happening is or a clear expected value for a given intervention, but they often give excellent answers to more comparative questions. "Does X seem like it would cost more per person than Y?" is easier than giving an exact number for either.

- Ensure that they have answered your question – If you ask questions such as “What are the main strengths and weaknesses of x intervention?” it is quite easy for them to forget the initial question once they have been talking about the strengths of the intervention for a few minutes. Follow up on this with something like “Thanks for outlining the strengths of x intervention; what do you think are the main weaknesses?”
- Give them space to think – Don’t move onto the next question immediately after they stop answering the previous one. Ensure that there is a small pause so that they can add something else if they think of it.

Recording

At the start of the interview, we ask the expert if they are comfortable having the interview recorded. This means that we don’t have to take notes during the interview so we can entirely focus on the questions and their answers to them, and it also makes summarizing the interview easier because we can go back and listen to their answer to x question again, for example.

To be safe, we usually record in two different ways.

1. Recording the interview through Skype. More information on this can be found [on the Skype website](#).
 - a. When recording through Skype, remember to click “Stop Recording” at the top of the screen before ending the call, otherwise there is a risk that the recording will not be saved.
 - b. Skype only keeps the recording for thirty days, so you should save the recording to your computer immediately. We also store these recordings in our Research Agenda -> <cause area> -> Expert Interviews folder.
2. Recording the audio of the interview using an external program.
 - a. We recommend [MP3 Skype Recorder](#) for Windows users and [Ecomm Call Recorder](#) for Mac users. Both can record automatically.

Reviewing as potential mentors

It is critical to use the research team’s extensive interactions with researchers to identify additional mentors. At the same time, the additional effort for a CE researcher should be limited. Hence, the review of researchers as potential mentors should be quick and fully integrated into the existing process that outlines conversations with researchers. Otherwise, the likelihood of poor retention or poor data quality is high.

Process

Brief survey: Each research conversation is reviewed with a very brief survey/review. This survey is fully integrated into the normal conversation-tracking process.

One-time assessment: In March/April 2020, or whenever suitable, the CE research and curriculum teams will reevaluate the initial reviews. Researchers with potential will be contacted appropriately and eventually sent the Mentor Application Form.

Brief survey

The survey is kept intentionally short to ensure consistent application. It will be automatically graded and return a suitability score.

Questions to ask

- 35 percent general cause neutral questions
- 45 percent customized cause (e.g. mental health)
- 20 percent customized person (e.g. expert X)

Cause-neutral questions

A little bit about the project: I am surveying <cause area> experts to get a sense of what would be the best areas to research and launch charity start-ups in.

1. Could I record our interview for the purpose of making more accurate conversation notes later?
2. Give an outline of the interview (types of questions, length, etc.)
3. Do you have any questions about the interview?
4. What got you interested in <cause area>?
5. How long have you been working in <cause area>?
6. All things considered (cost effectiveness, execution difficulty, what existing organizations are already doing, etc.) what specific organizations would you like to see founded in the next five years?
7. All things considered, what intervention or organization do most people think is effective but in your opinion is not? Why?
8. Are there any areas you think are neglected by current actors in the field? Why do you think these areas are neglected?
9. Are there any areas that seem unusually cost effective and evidence based relative to others in the sphere?

At the end:

1. What are good resources to read (blogs, books, podcasts)?
2. Do you know anyone who would be interested in talking to, mentoring, or supporting a new charity founded in this space?
3. Who else in the movement do you think would give valuable information about these sorts of questions?
4. Would it be possible for you to introduce me to them?
5. Do you have any questions for me?
6. This conversation was really helpful; would it be possible for me to write up a summary of some of the points we talked about and send you a copy for review? Some experts are happy for us to put up a published set of conversation notes on the topic. But we can also anonymize the conversation or combine it with points from other experts into a more overall view (eg., five out of twelve experts think that this is the most promising intervention) depending on what you are most comfortable with.

Mental health questions

Customized cause area questions:

There are some broad areas of mental health we are considering. For each area, we consider two questions.

- 1) How effective the area seems generally (for starting new charities) (below average – average – above average – the best intervention)
- 2) What might be the most promising specific things to do within an area. For example, one of the broad areas is therapy, and one specific area that some people think is promising is online apps for lower-income countries.

Areas

- Task shifting
- Peer support
- App-based therapy
- Skype-based therapy
- Direct therapy
- Therapy
- Social change
- Lifestyle change
- Screening
- Research
- Medication
 - Diet supplementation
- Medical procedures

- Government lobbying
 - Do you know of any existing organizations that are lobbying for subjective well-being/mental health?
- Corporate Campaigns
- Lifestyle Improvements (e.g. weighted blankets)

Do you think there are important broad areas that are not covered under one of these headings? Do you know of any broad resources that compare different global mental health interventions? How do you think different metrics of subjective wellbeing compare to DALYs? Which do you think is the strongest measure in the mental health space?

Customized person questions:

These will be customized questions to ask a specific person.

For example, a question to ask someone who mainly works in global health:

- What is your perception of the funding landscape for mental health compared to global poverty? Is a new organization likely to be constrained by funding, or would possible sources of funding have good counterfactuals?

Or asking someone with more knowledge about specific types of mental health interventions:

- How do you think task shifting compares to computer-based therapies in terms of cost effectiveness?

Animal questions

There are some broad areas of animal advocacy we are considering (see bulleted list below). For each area we consider two core questions: How effective does an area seem for starting new charities? and, What might be the most promising interventions within an area? (For example, within the area of food technology, one intervention that may be promising is lobbying governments to ensure fair labeling of plant- and cell-based products.)

- Corporate outreach
 - Meat reduction campaigns (eg. meatless Mondays in universities)
 - Chronic welfare improvements campaigns (eg. environmental conditions)
 - Acute welfare improvements (eg. slaughter)
- Governmental outreach
 - Welfare improvements lobbying (eg. environmental conditions); follow-up questions for welfare improvements:

- Acute (eg. slaughter, handling, transport) vs chronic (eg. environmental conditions)
- Increasing follow-through rates of existing campaigns vs new campaigns
- Positive welfare interventions (eg. environmental enrichment) vs negative welfare interventions (environmental conditions)
- Meat reduction lobbying (ban the advertising of meat, meat tax, ban imports from lower welfare countries etc.)
- Wild animal suffering (WAS)
 - Research
 - Targeted interventions
- Veg outreach
- Neglected areas e.g. crustacean interventions, farmed rodents, bugs for human use (entomophagy, snails, silkworms, insects used in research etc.)
- Research
 - Building the evidence base
 - Targeted research (eg. institutional ask)
- Food technology

Do you think there are important broad areas that are not covered under one of these headings? Do you know of any broad resources that compare different animal interventions?

Customized person questions:

These will be customized questions to ask a specific person.

For example, a question to ask someone who works with funders/works for a fund:

- Is there an intervention that the funders you work with are particularly excited about funding?
 - Is there some metric that these funders are always looking for in the interventions they fund (e.g., high cost-effectiveness vs. neglected areas)?

Or asking someone with knowledge of our priority countries:

- What are the top considerations and challenges we should keep in mind when working in <country>?

Family planning questions

Broad questions

1. How long have you been working in/researching family planning and what got you interested in it?

2. All things considered (cost-effectiveness, execution difficulty, what existing organizations are already doing, etc.), if a nonprofit was starting tomorrow, what specific program would you like to see implemented?
3. All things considered, what intervention or organization do most people think is effective but in your opinion is not? Why?

Comparative questions

1. There are some broad areas of family planning we are considering. For each area, we would like to get a sense of
 - a. How effective the area seems generally (for starting new nonprofits) (below average – average – above average – the best)
 - b. What might be the most and least promising specific things to do within an area? For example, you might think contraceptives supply and distribution is average compared to other broad areas, and within that community, health workers are most promising and social franchising is least promising.
 - c. Contraceptives provision/distribution/supply
 - d. Incentives
 - e. Information/education
 - f. Policy/advocacy
 - g. Research
 - h. Services delivery/ quality
 - i. Social and behavioral change
 - j. Women's and girls' empowerment
2. Which contraceptive is the most/least important to promote among, say, IUDs, pills, emergency contraception, condoms, implants, injectables, sterilization, and natural/traditional methods?
3. Do you know of any broad resources that compare different family planning interventions?

Funding

1. What kind of programs do you think are particularly easy to get funded, and which ones aren't?
2. Do you think there is funding available for new organizations to scale evidence-based programs?

Note-taking and summarizing

For every conversation, we ask the expert if we can take notes and share them, either named or anonymized as input from one of several experts we have interviewed on the topic. We offer to send them a copy of any notes we take so they can comment if they feel we misunderstood anything. We also offer to send them a copy of our full report if they are interested in seeing other experts' views or other synthesized research we conduct on this topic. If they are interested, we would love their feedback on the full report.

Conversation notes are summarized into an easily readable document and then sent back to the expert for confirmation that we did not misunderstand or misrepresent anything, similar to the [GiveWell conversation notes](#).

In an experiment using [Otter](#) (an automatic transcription service) to record the conversation vs. manual conversation notes (making notes whilst listening to the audio), we found that the time taken for the automatic transcription to be edited into a readable form is slightly longer than the time taken to manually write up conversation notes.

Synthesis

Although the bulk of the expert report will be the conversation notes, the project lead will have to synthesize these thoughts into a one-page easy-to-read summary. This can include both narrative explanation of concepts that came up multiple times in conversation notes or table-based data with rough quantification on what experts thought. An example of this can be seen in [this article](#). The section above the “additional details” is more reflective of the detail that would be helpful to have for a specific idea expert synthesis.

Expert review

Our final utilization of experts is when our full report is nearing conclusion. We send out our full report to any expert who indicated an interest in seeing the endline results and ask for any feedback they have.

Different levels of expert depth

Five minutes

Process

Speaking to an expert for five minutes is not possible because even finding and contacting them would take close to this amount of time. However, one broad expert conversation can cover a lot of ground. If there are 300 interventions in a given cause area and each is given five minutes, this adds up to twenty-five hours of total time. Of this, seven hours would need to be used to find and contact experts, two hours to prepare the most important questions that would give helpful information across an intervention area, ten hours to interview the top five who respond, and finally eight hours to synthesize the notes from speaking to them and scoring interventions based on the responses. This results in around five minutes per intervention in a given area or about five experts for the area.

Interviews at this level of depth will be used in our research agenda phase. We will contact five broad experts to help us narrow down our long list of ~300 <cause area> intervention ideas. For this level of depth, we will use an automatic transcription service (Otter seems best). We will send this transcription to the experts for review following our Skype call, explaining that this is an automatic transcription that has not been edited and will not be published.

Expected outcome

- Have spoken to several broad experts to get their sense of the area
- Summarized conversation and how it affects the intervention (one paragraph)
- Synthesized information from the interviews that will enable rating of each idea

Advice

- Contact at least three times the number of experts you want to talk to.
- Reach out to them before you start any of the methodologies. Even if EpV turns out to be the third or last methodology, they will take time to respond, so you can pursue the others in the meantime.
- Send your questions to them before your Skype call so they are aware of what kind of questions you will be asking and can prepare if they feel they need to.

Lessons learned

1. The categorization of interventions into a broad category (which you will ask experts to rate) is a highly important part of this process that we should have put more thought into for three main reasons.
2. It was hard to come up with broad categories that all interventions could fall under, but this was necessary because most experts would not have enough time to rate each individual intervention idea, so we had to get them to rate broad category areas instead.
3. Because all interventions did not fit neatly into the broad category areas, we sometimes had to make inferences about how they would score an intervention based on specific comments made during the interview.
4. In some cases, it was difficult to score individual interventions because they could fall under multiple broad category areas. For example, in animals, we had a lot of wild animal suffering interventions, so we came up with the broad category of “Wild Animal Suffering” (WAS) and got experts to rate this category. We also had a lot of interventions to prevent acute suffering, so we came up with the broad category of “Acute Suffering Interventions” and got experts to rate this category. In cases such as this, it was difficult to tell, for example, what category the intervention “using snap traps over glue traps for wild rodents” would best fall under: wild animal suffering, or acute suffering? There were a few instances like this that made some interventions difficult to score.

Potential solutions

1. Send experts the whole list of interventions and get them to rate each one individually, but this would be very time consuming for them.
2. Add subcategories (e.g., from the example above “Acute WAS Interventions” and “Chronic WAS Interventions”).
3. Send them a Google form survey that better illustrates the type of interventions that would fall under each category and get them to input their ratings there.
4. Ask experts about one representative intervention and infer from this the score for the rest of the interventions in that category/subcategory.
5. If we were to publish the results from the five-minute expert surveys, experts might be more interested in rating all of the individual interventions.
6. Remove the expert survey from the five-minute process as a standalone factor, but integrate it with informed considerations.

What we think would be best to do in the future: Incorporate both a Google form survey and the Skype interview into the EpV (e.g., Skype with experts and ask open-ended questions to get a sense of their values etc., and after the Skype send them a Google form survey and get them to rank subcategories of interventions on there, explaining what sort of intervention would fall into these categories).

Two hours

Process

Given the generalized information that has been gained from the five-minute process, the most helpful expert to speak to next would be a single domain expert. This interview would take one hour including prep, contacting the expert, and summarizing the notes afterward, as well as one hour speaking directly to the expert about the key questions that would be hard for a broader expert to answer. If an expert can cover more than one cause area, more time can be used to prep for his or her interview. Over a given cause area this would result in one expert per domain or five to thirty experts in total. This stage will only happen for family planning this year. In other cause areas, we will use a different two-hour methodology.

Interviews at this level of depth will be used to help us understand a more specific charity idea such as what a promising country or approach might look like. These interviews, which will often be with domain experts, will help us determine, for example, what country would be most promising to run a conditional cash transfer (CCT) program for intrauterine devices (IUDs) in. For this level of depth, we will manually write up conversation notes (because these notes will be published, they need to be higher quality than for the other levels of depth). We will send this conversation summary to the experts for review before publishing. We will also invite these specialist experts to review the report as a whole.

Expected outcome

- Have spoken to at least one domain expert to get their sense of the area
- Summarized conversation and how it affects the intervention (one paragraph)
- Synthesized information from the interviews that will enable rating of each idea

Advice

At this stage, you could offer the following options to the experts you are interviewing:

- You can send your questions in a Google document and they can write their answers on there, with a quick follow-up Skype to ask any further questions if necessary
- One-hour Skype interview
 - If they choose this, you should send your questions to them before your call so they are aware of what kind of questions you will be asking and can prepare if they feel they need to.

Twenty hours

Process

At the twenty-hour level, roughly the same process would be used as for the five-minute level. Of this time, two hours would need to be used to find and contact experts, as well as two hours preparing the most important questions that would give helpful information across an intervention area, six hours interviewing the top five that respond, four hours synthesizing the notes from speaking to them and scoring interventions based on the responses, and six hours sending the results to the experts and getting their feedback on both their personal notes and the report as a whole. This would result in five experts spoken to about a single intervention.

These expert interviews will often be with specialist experts, such as speaking to an animal advocate in Taiwan or speaking to a fish disease specialist to help us determine whether paying farmers to use vaccines to treat diseased tilapia in Taiwan seems like a promising intervention for farmed fish. One of the experts contacted will be a person leading an organization implementing similar interventions in the same region as the recommended charity. The goal of this interview is to find out if they could be influenced to change their program to a more cost-effective one (recommended by CE). Such a change would alter where we plan to allocate resources. For example, influencing a charity fortifying flour with iron to add folic acid would alter the score of a folic acid fortification intervention, and would lead to starting a tobacco taxation charity instead.

For this level of depth, we will manually write up conversation notes (because these notes will be published, they need to be higher quality than for the other levels of depth). We will send this conversation summary to the experts for review before publishing. We will also invite these specialist experts to review the report as a whole.

Timeline relative to other methods (eighty-hour report)

Ten hours – Broad undirected reading and crucial consideration, CC (IC)

Sixteen hours – Directed research (WFM)

Ten hours – Finding and talking to experts (EpV)

Twenty hours – CEA creation (CEA)

Four hours – Directed research (WFM)

Ten hours – Summary writing and internal contemplation (IC)

Ten hours – Showing endline report to experts (EpV)

Expected outcome

- One-page summary of the synthesis expert conversion
- Conversion notes for all experts interviewed
- Expert feedback given on the final report
- An expert or two highlighted who would be willing to mentor or speak to a new charity

Advice

At this stage, you could offer the following options to the experts you are interviewing:

- You can send your questions in a Google document and they can write their answers on there, with a quick follow-up Skype to ask any further questions if necessary.
- One-hour Skype interview
 - If they choose this, you should send your questions to them before your Skype so they are aware of what kind of questions you will be asking and can prepare if they feel they need to.

External expert data

Experts we interview are not the only source of expert data we use. If there are previously written interviews, conversation notes, or other direct sources of expert data, we also include these in our expert report. Data like these would be searched for in the directed research phase of the project but would be included in the evidence section of data.

Summary

By the time a charity is recommended in an area, we will have spoken to five broad experts in the cause area, as well as six domain experts or specialist experts. Some of these experts will also have reviewed the overall report and given comments or suggested improvements. We also will have taken into account any publicly available expert surveys or summaries of other related conversation notes.

These experts are spoken to using a consistent methodology. Our conversation notes, as well as our summarized interpretation of the conversations, are published in a single section but clearly differentiated from each other.

Reading list

- [External resources on how to be a good interviewer](#)
- [Evidence on good forecasting practices from the Good Judgment Project: An accompanying blog post](#)
- [The Black Swan: The Impact of the Highly Improbable](#)
- [Expert Political Judgment: How Good Is It? How Can We Know?](#)
- [Future Babble: Why Expert Predictions Fail and Why We Believe Them Anyway](#)
- [Scientists can sense which studies are weak](#)

2. How Do You Know Who to Trust?

Author:

Joey Savoie,
Co-founder & Director of Strategy

When you are wheeled into a hospital with a broken arm, you place your trust in multiple people and establishments. You trust that the nurses are giving you the right medication to prepare you, and that the doctor will make the right call on how to fix your arm. And yet you personally know relatively little about the specific treatments that are going to be applied.

We defer to people all the time on different issues, whether it's the doctor at a hospital, the weatherman for the forecast, or the baker who tells us the bread is fresh. Even in our domains of expertise many judgment calls are made by others, and we have to trust or distrust their data.

Knowing who to trust is both a difficult and important skill. Trust the wrong person, and they can fill your head with wrong information. But trust no one, and you have to fix every broken bone yourself. So how can we determine who to trust – who is credible and who is not?

There are four key ways to determine whether a source or person is worth putting your trust in. In descending order of how good an indicator it is, you can:

- Check against reality
- Check against further research/reasoning
- Check the source's reliability in other areas
- Check for signs of credibility signaling

Each of these is more of a spot check than perfectly predictive, and not all can be done in every case.

Check against reality

The best way to test if you can trust a person or source is to check their statements against reality. Say there are two weathermen and you are unsure who to trust. In this case, a reality check is easy. You could compare each of their historical predictions with the historical weather to see which has been accurate more often. This does not guarantee who will be a better source in the future, but it's strongly suggestive. Similarly if a source predicts a certain reality, particularly in a manner that is easily **falsifiable**, this evidence can be used to support or create skepticism for the sources.

Reality checks can also be used for groups of sources. For example, lots of people go to the hospital with broken arms and generally come out with a cast and an improved state. Thus I might generally trust hospitals to fix broken arms, even if I have not checked my specific doctor. **Reality is the ultimate arbiter**. It does not matter if one weatherman speaks more confidently, wears a better suit, or has a PhD – the one whose predictions more closely correlate with reality is the better source.

Check against further research/reasoning

Not all claims can be checked against reality, but a large number can be checked with further research. For example, the first time I heard from GiveWell that global poverty was reducing over time I was surprised. However, when I checked multiple other sources it indeed looked like this was the case. This made me trust GiveWell's research more. If it had turned out that global poverty was in fact on the rise, I would have been more skeptical of their research.

Enough spot checks and slowly a source as a whole can become trustworthy. For example, I have checked over a dozen different sections of GiveWell's work, often putting in several dozen hours of research into a specific claim. Again and again, from my best assessment it looks like they are correct. Over time this builds trust, so that now I can use GiveWell as a reliable source to check other claims against.

Check the source's reliability in other areas

Trustworthy in one domain does not always mean trustworthy in another. Despite the hospital fixing my arm, I would be wary of their ability to predict the weather.

However, sources can often make claims in some areas that are testable and others that are not. In this case it can be useful to look into a source's reliability in a different area. For example, imagine I have a friend who for fun memorizes facts that my research shows consistently to be correct. Were this friend to share a new fact in a hard-to-check area, I would be likely to believe him. Likewise if a source is highly trustworthy in one area, it's more likely to be trustworthy in others. If GiveWell started recommending animal charities, I would be inclined to trust these recommendations even if I had not yet checked them against reality or further research.

Check for signs of credibility signaling

The last way you can try to get a sense of who to trust is by looking at generally accepted forms of credibility signaling (e.g. legitimacy of the source, or an individual's qualifications). This has the advantage of being quick, but it's also fairly unreliable compared to the other methodologies. A nice website signals that a source has funding, but is a pretty weak signal in terms of them being trustworthy. Credibility signaling is often where people go wrong with trusting a source – by giving a certain signal far too much weight compared to its actual correlation with reality.

Summary

When assessing a person or piece of evidence, we often look at multiple sources. If I am assessing the credibility of a study, for example, I might first look for replication studies (checking against **reality**). If those cannot be found, I would check the methodology of the original study (checking against **reasoning/process**) to see if it had been well run and seems internally valid. If that doesn't work I would look for other studies and sources that point in the same or opposite directions (checking against **further research**). If I find some but not enough to be convinced, I might look at the study authors' other studies or the work they have done in an area I know quite well (checking against **other areas**). Finally, I might look at what school the authors were from or what journal they published in (check for **credibility signaling**).

When any of these sources conflict, it's easy to see which one wins. **I care far more about the source that is consistently in line with reality than I do about the source with the most PhDs.**

Over time as sources become highly reliable, you can use them to check other sources. If you check **Cochrane** a dozen times and it dovetails with further research, you can eventually start to use it as a reliable source to check others against. Over time your trust can become more and more refined. For example, you might only trust certain sources in some areas and not others.

XI. Budgeting & financial planning (the saw)

The finest wood projects require as many cuts as they require additions. When making decisions, you will have to determine how to plan around limited resources such as time and funding. Knowing what areas to spend less on or to cut altogether is an important skill often learned later in one's decision-making journey. Like the saw itself, if misused it is easy to cause injury to yourself and others and it is a scary tool for those new to using it. However, it is as essential as the hammer for any great project. For every choice, you have to say no to hundreds of other paths, directions, and opportunities. In this section, we will cover:

- How to know which projects to cut
- Reasonable safety nets for an entrepreneur

1. Five Ways to Cut Costs

Author:

Patrick Stadler,
Charity Mentor & Director of Communications

The fundraising pitch went well, and thirty days later, you shout in excitement as a six-figure number pops up on your e-banking app. But don't get too pumped. Now begins the quest of spending your resources diligently. Here are a few key guidelines to consider:

1. Spend only on one or two high-impact projects
2. Cut fixed costs
3. Use an activity budget
4. Start with a low salary
5. Watch out for common cost drivers

Spend only on one to two high-impact projects

As a curious founder, you see so many opportunities for your charity. It is tempting to pursue many of them at the same time. Yet that would be a mistake. By spending on various projects, you allocate funding to opportunities below the top choice in terms of impact. Why spend on the second, third, and fourth most cost-effective project if you can instead pool resources to the best choice?

As you spend on more than one project, you are not just reducing your monetary power but also your managerial focus. You already have limited bandwidth with one project as a charity founder. Don't stretch it further. So as a general rule, focus on one project at a time – or in exceptional circumstances on two.

Cut fixed costs

A trendy office in the heart of London: doesn't this qualify as an investment in your staff and your organization's long-term success? Watch out before you sign that two-year lease in the glass skyscraper. While possibly obvious in this example, the choice for which fixed costs make sense is not always easy to spot. The problem with such costs is that they increase your overhead while not having a direct relationship with the number of beneficiaries you serve. The size and interior of the office do not influence your potential to fulfill your charity's mission. Many organizations, including charities such as New Incentives, operate completely **remotely**. Others like Charity Entrepreneurship have successfully found affordable offices even in global cities such as London.

Reducing fixed costs applies to staff as well, most likely your number one expense item. Here it is less straightforward as you need skilled staff members to serve your beneficiaries. You also don't want to burn out or stretch your managerial capacity too much by not recruiting. Yet as an early organization, stick to the co-founding team for as long as possible. Your program might still iterate quite a bit and new hires might not fit into an organization that has undergone a pivot. While staff increases capacity, it also requires more oversight on your end. Additionally, being a smaller team forces you to focus on your strategic decisions. And of course, fewer staff equals lower cost.

Use an activity budget

Your overall budget includes all expenses from program to overhead. Your activity budget, on the other hand, tracks all costs going into serving one beneficiary. As such, it allows you to track your unit costs. Working with unit costs helps you design a cost-effective and scalable program. If your beneficiaries are spread out over hundreds of villages or clinics, for example, you will have high transport costs. Noticing this in an activity budget will start the brainstorming process on how to serve beneficiaries in a more sustainable way (e.g. by organizing camps that bring beneficiaries together in fewer locations).

Start with a low salary

We sometimes get questions from applicants whether they could assign themselves an annual salary of, say, \$75,000. Technically, they can. After all, you call the shots as founder. Yet this might not be the best idea. A relatively high salary at a start-up has several downsides:

- It limits your ability to pivot as your runway is shorter (all things being equal).
- It makes fundraising harder by requiring you to potentially raise double or triple of what you would have asked for otherwise.
- Depending on your donor, it can be seen as too high for an organization with an unproven track record.

Therefore, we encourage our incubatees to start out with relatively low salaries. As an organization we also believe in **a culture that does not align seniority with a higher salary**. Our co-founders, for example, are among the lowest-paid staff members. We encourage our charities to follow suit and actively decouple personal worth as an employee from the monthly paycheck.

Watch out for common cost drivers

Aside from salary and office, there are several other common cost drivers at charities.

- Field transport can be a crucial one. A developing country context does not mean that transport will be particularly cheap, especially if you have many staff members who need to travel on a frequent basis.
- Electricity is another, as the grid can be unreliable in many countries and running a generator is expensive. As a modern organization you likely run

your program digitally, so you need to factor in the costs for getting each staff member at least a mobile phone and accessories such as power banks.

- Finally, training costs can be higher than expected – as you usually underestimate the duration needed to onboard new field staff and facilities such as training halls are costly.

In sum, there are many ways to make the best of your charity's finances as a founder – so make sure to protect this six-figure number on your start-up's bank account for as long as possible...

2. What Is a Reasonable Financial Safety Net for a Charity Entrepreneur?

Author:

Joey Savoie,
Co-founder & Director of Strategy

Charity entrepreneurship is in many ways a less stable career than a traditional job. Charities in their early days will have limited runways (often under six months). At any point, results could come in showing that an intervention is not worth continuing. On the other hand, value drift is an important consideration. Losing motivation to achieve altruistic impact happens more quickly than one would expect (see [Empirical Data on Value Drift](#) and [Concrete Ways to Reduce Value Drift](#)). Hence, it is better to donate more to effective causes now than to accumulate a large safety buffer that you might spend ineffectively in the case of value drift. So a question arises: how much personal runway or savings should someone have when becoming a charity entrepreneur?

In the case of any job, people will want enough safety nets to be able to securely transition from one job to another. The first fact to consider is the **average length of time between jobs**. It's worth looking at both your past history and the history of similar peers. The second thing to consider is cash-based savings, but there are also many often forgotten safety nets. Monetary savings are only one of them – frequently representing a fairly small percentage of a person's overall safety nets.

Often when non-financial safety nets are taken into account, everyone's position looks substantially more secure, and thus more confident. **Non-financial safety nets are things that exist in a peer group or world outside of a personal bank account, that extend the effective time a person could be unemployed for without considerable negative ramifications.** Living on a friend's couch for months wouldn't be that thrilling, but it would be far from homelessness.

Some factors to consider when estimating your current personal runway include:

Partner

For those who have a partner, this will often be the first level of safety net. Even if a major event throws your earning ability out the window, often a single earner (particularly in high-earning demographics) can cover the difference.

Cost-cutting

This is a strategy that likely would work in combination with any others. Few people are cutting expenses as much as they could if the need arises. Certain factors like rent might be hard to quickly change, but food and entertainment budgets are generally a lot more flexible.

Family

Not all families are able or willing to provide long-term support, but many are. Some people, due to having a partner or divorced parents, have multiple plausible families who could help out with things like free rent, and in many crisis situations, most would be more than happy to help for fairly long periods of time.

Internal capacity

Some events requiring safety nets might completely destroy one's effective capacity to work. However, many would just lower the value of the job I would get if I had to find work in a rush due to financial pressure. For example, many people skilled enough to be charity entrepreneurs could get an entry-level job in some industry within a matter of weeks if they needed to.

Friends

Not everyone has a group of friends that would be able to give a temporary place to live, or friends wealthy enough to give a loan, but many people do. People from the charity sector, in particular, tend to have more altruistic and high-talent peer groups. In many cases of serious distress, friends can pull a lot of weight.

Possession-based safety nets

Many people own a house, a car, or other assets that can be used as an extra safety net, although not a quickly accessible one, in times of need.

Legal insurances

Many individuals have several insurances such as home, car, and medical. The point of these insurances is to cover unexpected negative events, although they are category-specific.

Government safety net

Though not the most prestigious option, many government programs are available for people who experience a sudden shock. The details vary a lot from country to country, but social programs such as employment insurance, welfare, and subsidized living exist in many high-income countries.

Credit and loans

40% of US households carry credit card debt. Given the high interest rate of 12-15%, spending on credit should only be used as a last resort. Cheaper options for loans include crowdlending online platforms and mortgage providers (if you own a house). In general, debt should be avoided if at all possible. If necessary, look into no- or low-interest loans from family and friends first (see above).

Combine safety nets to get optimal protection

One can imagine combining these resources in different ways and ending up at **12-24 months of safety net** or more while having nothing in cash savings. These were the ones that came to my mind. Everyone's personal situation will be different, but I think that listing these considerations can generally reduce anxiety around savings buffers.

I expect individual charity entrepreneurs who apply consideration to their full scope of resources would generally feel more confident donating or taking on ambitious projects without large savings buffers.

XII. Cost-effectiveness analysis (the magnifying glass)

Sometimes a deep look is needed to understand the details of something. In the world of woodwork, you might use a magnifying glass to get a closer look at an aspect of the project. When it comes to entrepreneurial projects a closer look is achieved through a formal cost-effective analysis (CEA). This closer look is crucial for understanding a project and evaluating it in-depth and involves putting all the variables into a single complex equation. In this section, we cover:

- What cost-effectiveness analysis is and why it's important
- Metrics commonly used to quantify complex moral values in CEAs
- Easy mistakes to make and how to avoid them
- How and when to use speculative and rigorous CEAs
- How CE uses cost-effectiveness analysis

1. Introduction to Cost-effectiveness Analysis

A cost-effectiveness analysis (sometimes called “cost-benefit analysis” or “cost-utility analysis”) tries to come up with a single metric to compare the benefits of two or more different interventions side by side.

Imagine **two charities**, one of which can train one guide dog to assist one blind person for each additional \$1,000 donated, and one of which can perform one vision-restoring cataract surgery for each additional \$35 donated. While performing cataract surgeries and training guide dogs are two very different interventions, once you recognize that both of these interventions are fundamentally aimed at ameliorating blindness, you can confidently say that performing cataract surgeries is a more cost-effective way to achieve that goal.

When comparing cataracts and guide dogs, the appropriate metrics would all have to do with vision and blindness. The specifics of this can get quite complex. For example, “number of blind people helped” would be the simplest metric – you

could probably just add up the number of dogs and the number of cataract surgeries. However, you might think that addressing 20 years of blindness is better than addressing 10 years of blindness, and choose “years of blindness ameliorated” as a metric instead. Now, you will have to think about many other considerations, such as how many years after deployment the average guide dog can productively work, how long the average cataract surgery delays blindness via some other means, how long the average beneficiary will live anyway, and so on.

When comparing every intervention in the world, even within a single cause area, things get far more complex. Is preventing a year of blindness worth missing the opportunity to prevent a year of depression? What if the blindness also reduces income, or causes depression? Is the quality of life improvement from doubling the income of 50 impoverished people worth missing the opportunity to save the life of one young adult who would have lived for 50 more years? What if that young adult was also supporting a family containing several dependents?

As you can see, these are not questions with easy answers. Figuring out how to count these things involves not only philosophical uncertainties about what is good and empirical uncertainties about what is true, but also basic uncertainties about how to count things and the potential for extremely different answers based on minor differences in methodology. How you answer these questions can *drastically* shift which charities you consider to be the most cost-effective. Things, of course, get even more complex once you bring in concepts such as improving the experience of animals, or steering the trajectory of the future.

The next few articles will go over some useful methodologies for making these judgment calls and calculating the cost-effectiveness of an intervention. It will also explain how these calculations can be used to determine the distribution of funding, calculate your impact, and make decisions about your charity.

Resources

[Quantifying Complex Values](#): QALYs, DALYs, moral weights, welfare points, and subjective well-being

[13 Common CEA Mistakes](#)

[Using Speculative CEAs to Make Decisions](#)

[Using More Rigorous CEAs: New Incentives, Fortify Health, and CSH](#)

Other resources

[How CE Uses Cost-effectiveness Analysis](#)

GiveWell: Cost-Effectiveness Overview

Economic Evaluation Webcast Part 5 of 5: Cost-Effectiveness Analysis

Reference Case Guidelines for Benefit-Cost Analysis in Global Health and Development

2. Quantifying Complex Values

Author:

Ishaan Guptasarma,
Curriculum Developer & Content Writer

There are many different ways to measure positive outcomes. Depending on which metrics one chooses, one might end up with very different cost-effectiveness estimates for an intervention. Here are some of the common metrics we used and encountered in the course of our research process.

QALYs and DALYs

QALYs (quality-adjusted life years) and DALYs (disability-adjusted life years) are commonly used metrics that composite multiple measures of health and well-being to arrive at a single number.

One QALY is one year of perfect health. A year spent bedridden might be represented as 0.5 QALYs, while death would be represented as 0 QALYs. A state worse than death might be assigned negative QALYs. For example, one might imagine that a terminally ill patient experiencing extreme and irreversible cognitive decline, depression, and high amounts of pain falls into this category.

Health states are assigned values between 1 and 0 (or lower) within the QALY system by asking individuals various hypothetical questions. These hypothetical questions include time trade-off questions (how many years of life would you trade to be cured of this condition) and standard gamble questions (would you undergo a procedure to be completely cured of this condition if it had an x probability of killing you?).

One of the weaknesses of QALYs is that these questions are generally asked of the general population, rather than people actually experiencing the problems in question. This leads to some paradoxical results. For example, many people who are assigned negative QALYs by one methodology (a state supposedly worse than death) **report being happy and satisfied with their lives** when questioned under a different methodology! This subjectivity and lack of sensitivity make it especially difficult to

apply QALYs to contexts where respondents may not have sufficient experience to give accurate weights, such as with chronic illnesses.

DALYs (Disability Adjusted Life Years) are quite similar to QALYs. The main differences are that DALY calculations sometimes explicitly include age-weightings (so extending the life of a young adult by one year might be more important than extending the life of an elderly person by the same amount), and that weightings for various conditions are generally chosen by experts, rather than elicited by presenting hypothetical choices to a large sample of people.

A DALY is the sum of two components: *Years of Life Lost (YLL)* and *Years Lived in Disability and Disease (YLD)*. If you died at age 50 but your life expectancy was 82.5, you would have suffered $82.5 - 50 = 32.5$ YLL. If you also went blind at 45, you would multiply the 5 years spent in blindness by some weighted factor. If you think that giving someone 1 year of additional life is as good as giving someone 3 years of additional sight, you would assign this weighting to be $\frac{1}{3} = 0.33$ to get $0.33 * 5 = 1.65$ YLD. According to this calculation, averting 5 years of blindness is as good as giving someone 1.65 years of healthy life. In total, you would have suffered $32.5 + 1.65 = 34.15$ DALYs.

To summarize:

$\text{DALY} = (\text{life expectancy} - \text{age of death}) + (\text{years spent disabled} * \text{weighted factor})$

$\text{DALY} = \text{YLL} + \text{YLD}$

In the example above, that's $34.15 = 82.5 - 50 + 5 * 0.33$

Some DALY formulations will additionally weight some years as more valuable than others, such that years spent as a young adult are given more value than years spent as a baby or as an elderly person. Even without such weighting, DALY methodologies will tend to direct priorities toward saving the lives of the young and able-bodied, raising various ethical dilemmas.

QALYs, DALYs, and similar metrics are sometimes collectively referred to as HALYs (health-adjusted life years). They are used by the British National Health Service, the World Health Organization, and many others to make recommendations and policy decisions. See [Sassi, F. \(2006\)](#), and [Measuring burden of disease: the concept of QALYs and DALYs](#) for more details about calculating DALYs and QALYs, and see [GiveWell's blog on DALYs](#) for more discussion on their weaknesses and why not to use them.

GiveWell moral weights

GiveWell **does not exclusively use DALYs** because DALYs combine various outcomes into a single metric, and different stakeholders **disagree** about which outcomes are good. Instead, GiveWell staff assign **moral weights** to various outcomes, such as preventing child death, preventing adult death, doubling income, or improving health. These weights are partly **informed (but not fully determined) by the judgments** of populations who are demographically similar to beneficiaries.

Many **GiveWell CEAs** include a **section** in which stakeholders are encouraged to make a copy of the spreadsheet and input their own moral weights, to see how different value judgments shift which interventions become more cost-effective to work on or donate to.

GiveWell moral weights are combined with real world data to determine how cost-effective different interventions may be relative to each other.

Consideration	AMF	Malaria Consortium	Deworm the World	END Fund	SCI	Sightsavers	GiveDirectly
Estimated cost-effectiveness (relative to cash transfers)	~4x	~4x	~10x	~4x	~8x	~5x	Baseline
Primary benefits of the intervention	Under-5 deaths averted and possible increased income in adulthood		Possible increased income in adulthood				Immediate increase in consumption and assets

Table source: *Updated top charities for giving season 2016*

Charity Entrepreneurship's Welfare Points

When we were researching cost-effective interventions into the animal space, we felt that there were no good cross-species metrics available for comparing welfare. While there are methods used by researchers, activists, and industry for assessing health and well-being within a single species, there was no good way to compare interventions targeted at factory-farmed cows with interventions targeted at wild rats. To remedy this, we created **Welfare Points**. Welfare points combine a range of metrics that apply across species, including rates of premature mortality, malnutrition, and injury. We used welfare points as a tool to recommend interventions for this program.


Animal	Overall	Welfare score
Max score	Priority	100 to -100
Wild fish for human use	Very High Priority	-47
Wild fish	Very High Priority	-31
FF Fish-traditional aquaculture	Very High Priority	-44
FF Turkey	High Priority	-57
Wild bug	High Priority	-42
FF Broiler chicken	High Priority	-56
Wild rat	High Priority	-28
EU FF laying hens (enriched cages)	Mid Priority	-46
Wild bird	Mid Priority	-2
USA FF layings hens (battery cages)	Mid Priority	-57
Wild Chimp	Low Priority	47
Human	Low Priority	32-81
FF Cow milk	Low Priority	-34
FF Beef Cow	Low Priority	-20

Image: CE *Weighted Animal Welfare Index*

Subjective well-being and life satisfaction

While QALYs, DALYs, and GiveWell moral weights generally focus on outcomes such as preventing death and disability or increasing income, Happier Lives Institute makes the case that **income and health outcomes don't necessarily reflect happiness**. If someone valued happiness more than other outcomes, they might choose to work on or donate to interventions directed at maximizing metrics that aim to measure happiness more directly.

Some **methods to measure happiness** include experience sampling methods (asking people how they feel at random times during a day), day reconstruction methods (asking people to rate the previous day), and life satisfaction data (asking people to rate their life satisfaction out of ten).



*Image: How QALY time trade-off assessments differ from day reconstruction and life satisfaction assessments (Image from HLI on **Health metrics**, Data from **Dolan and Metcalfe 2012**)*

There are various reasons to be skeptical with data regarding subjective well-being.

- People are not consistent in their responses, giving different answers to different methodologies.
- There is a tendency for most people to say that they experience above average happiness, regardless of circumstance (similar to how most people self report that they are **above-average drivers**).
- There are extreme cultural differences. Not only are there starkly different cultural conceptions of happiness, but even the concept of rating something out of 10 can be difficult to explain in some settings.
- People in difficult circumstances assign higher numbers to worse states, while people in easier circumstances assign lower numbers to better states, such that one person's 6 is another person's 9.

- Questionnaires that aim to calibrate these responses to one another are longer and more difficult to explain.
- It's difficult to know how these numbers might tie into underlying morally relevant conditions.

Despite these drawbacks, subjective well-being ratings have the advantage of accessing data directly from beneficiaries, unlike the other metrics that have been discussed. They also are appealing from the perspective of moral intuitions. Many believe that happiness and life satisfaction are ultimately what matter, and there is a convincing case that other commonly used metrics don't function as good proxies for these outcomes.

Discounting and adjustments

Two interventions might show a similar number of QALYs, welfare points, lives saved, etc. on paper, but might in practice be supported by different levels of evidence, occur over different time frames, or have other extenuating factors that change one's view of their true cost-effectiveness. Applying **discounting factors** is one way to address these issues.

Time discounting

In an uncertain world, you ought to prefer receiving \$100 today to receiving \$101 next year. This is true for a variety of reasons: you cannot be certain that you won't need the money this year; you may not be certain that the person will keep their promise to give you the money next year; and besides, you can probably invest \$100 today at an interest rate greater than 1% and come out ahead next year anyway. For the same reason, you ought to consider immediate good outcomes as better than identical good outcomes projected to happen in the uncertain future. For example, if one of the benefits of your deworming intervention is that it **increases income later in life**, you should add a time discounting factor.

(The benefits of having an impact sooner is also one of the arguments for **donating money sooner rather than investing to donate later**, although for that question there are **considerations to balance on either side**, such as the possibility of new donation opportunities appearing in the future, or of improving conditions making helping people more expensive over time).

Evidence discounting

It's generally observed that the more certainty one acquires about a cost-effectiveness estimate, the less optimistic the estimate becomes. There's many reasons for this. One such reason is regression to the mean. In general, most interventions do not surpass strict thresholds of cost-effectiveness. Whenever you happen upon something that *seems* cost-effective, chances are that this is due to a problem with your evidence or your model. This is analogous to publication bias in the sciences – exciting results are less likely to be true. Therefore, it's generally wise to apply a discount factor to an intervention that you have less evidence about.

GiveWell's [2019 GiveWell Cost-Effectiveness Analysis](#) includes a functionality where a stakeholder can enter a [time discount](#) factor, to represent a preference for good to happen now rather than later. There is also a [replicability adjustment](#) regarding deworming, to adjust for the fact that cost-effectiveness estimates for some deworming interventions were done by generalizing results from a study which occurred under some unusual conditions (El Niño-associated flooding occurred during the study, which may have altered epidemiological factors regarding schistosomiasis).

Summary

There are many ways to quantify good, including QALYs and DALYs, GiveWell moral weights, Welfare Points, and various subjective well-being and life satisfaction measures. Each metric comes with a set of implicit value judgments and methodological limitations. Such metrics do not always account for various real world factors regarding evidence and timing, so you must account for such considerations by applying adjustments and discount factors. GiveWell spreadsheets often attempt to separate subjective judgment calls from concrete evidence, and try to design their cost-effectiveness analyses such that stakeholders can adjust these values to reflect their own judgments. Presenting subjective calls separately from concrete evidence in this way is generally a good practice.

Resources

QALYs/DALYs

- [Sassi, F. \(2006\). Calculating QALYs, comparing QALY and DALY calculations. Health policy and planning, 21\(5\), 402-408.](#)

- [GiveWell: Disability-Adjusted Life Years: Introduction](#)
- [Measuring burden of disease: the concept of QALYs and DALYs:](#)

GiveWell moral weights

- [Interpreting the Disability-Adjusted Life-Year \(DALY\) Metric](#)
- [DALYs and disagreement](#)
- [New research on moral weights](#)
- [How GiveWell and mainstream policy makers compare the "good" achieved by different programs](#)
- [Approaches to Moral Weights: How GiveWell Compares to Other Actors](#)
- [GiveWell's Cost-Effectiveness Analyses](#)

Welfare Points

- [Is it better to be a wild rat or a factory farmed cow? A systematic method for comparing animal welfare.](#)

Mental health

- [Measuring happiness \(Post series\)](#)
- [A Happiness Manifesto: Why and How Effective Altruism Should Rethink its Approach to Maximising Human Welfare](#)
- [Cause profile: Mental health](#)
- [What Should A Utilitarian Billionaire Do To Maximise Happiness?](#)

3. 13 Common CEA Mistakes

Author:

Ishaan Guptasarma,
Curriculum Developer & Content Writer

Cost-effectiveness analyses are not necessarily intuitive. Over time, we've noticed some mistakes that tend to be made repeatedly. We've tried to list some of the most common ones here.

Taking cost-effectiveness estimates literally

If you've been following along with the rest of the chapters, you will have noticed that cost-effectiveness analyses involve many judgment calls – not only philosophical matters regarding morality and epistemology, but also much more arbitrary decisions about how to count things and which particular equation to use in a given scenario. In reading about this, we hope you've developed an intuitive sense for why you **can't take expected value estimates literally (even when they're unbiased)**.

Comparing dissimilar CEAs and failing to apply equal rigor

Because of all the aforementioned arbitrary decisions involved in creating a CEA, the precise numbers you get are often an artifact of the methods you used to get them. However, if two interventions were evaluated using the *same* CEA methodology, using similar models which operate under comparable assumptions and judgment calls, they can be directly compared in a productive manner.

If one were to compare two completely different CEAs that were constructed under completely different methodologies, it's quite likely that the differences in the final numbers are mostly a result of methodological artifacts, not real differences in the impact that the interventions have on the real world.

If you want to compare your CEA with a CEA created by GiveWell or another organization, you ought to model your own CEA to be sufficiently similar. **Equal application of rigor** is essential when comparing two interventions.

Committing the 1% fallacy

The **1% fallacy** is a phenomenon in which entrepreneurs pitch investors on a big, speculative idea, and then claim that even if they could only capture 1% of the market share or have a 1% probability of success, it would still be a good investment. Investors know not to fall for this pitch, because “to capture 1% of the market share” is actually an ambitiously large claim. “1%” is often a lot bolder and less reasonable than it seems.

Cost-effectiveness analyses for interventions that are difficult to measure and aim to have a massive effect (for example, policy change or corporate campaigns) can be vulnerable to this issue. The problem with the “even if we discount this by 99%” defense is that in real life, realistic probabilities can be and often are much lower than 1%.

Simplistic discounting/probabilities

Proposing to accomplish something which has enormous benefits and then simply adding “I conservatively assume a 1% chance of success in this endeavor” will lead to incredibly optimistic estimates.

Discounting and incorporating probabilities must be done separately for every assumption in your process, not just tacked on to the end of your analysis. If your intervention relies on ten separate assumptions to be true, and each of those assumptions comes with a 50% discount, the cumulative discount is actually 0.098%, an order of magnitude less than 1%.

Simplistic worst-case scenarios

Your analysis rests on knowing the values of several difficult to estimate quantities (e.g. the efficacy of an untested antidepressant, the number of crustaceans that exist, or the externalities of an unprecedented policy change). Suppose you tried doing the cost-effectiveness analysis under worst-case assumptions for these unknown values, where for each assumption you are 95% certain that reality is more favorable than this.

What's wrong with this? The more assumptions you add, the more hopelessly optimistic your so-called worst-case analysis becomes. If you have five worst-case assumptions where you are 95% certain that reality is more favorable than this, there's actually only a $95\%^5 = 77\%$ chance that one of the factors isn't catastrophically worse than you thought it was – that's hardly a worst case. Any average-case or best-case analysis using these methods runs into similar issues.

When dealing with a probabilistic range of values, put it into **Guesstimate**. Guesstimate is a tool that allows you to describe the probability distribution of possible values, and then runs a Monte Carlo simulation to find the range of likely outcomes.

Meta traps, double counting, and taking over the world

Meta-charity can dominate direct charity in terms of cost-effectiveness analysis.

For example, so long as you are raising more money than you are spending, fundraising for a high-impact organization will seem more cost-effective on paper than actually *being* a high-impact organization. (In reality, in order to actually understand whether fundraising is a high-impact opportunity, you'd have to factor in the opportunity costs – the counterfactuals of what the staff who worked on fundraising might have otherwise done with their time.) In addition to money, this logic can be applied to talent. Rather than working on an issue directly, it might seem more effective to recruit two or more staff of similar talent to work on it. What if we recruited two or more people *to work on recruiting*? Why, we could spread effective altruism everywhere, eventually taking over the world!

There is a degree to which this sort of thing does actually work. However, overdoing this type of thinking is what we call a **meta trap**. In reality, not everything depends on us – the people we recruited might have been recruited by another high-impact organization, or we might have figured out another way to be high-impact without recruiting them. The charities we funded might have figured out how to fundraise or otherwise realize their impact in a different way. The donors we redirected might have been redirected a different way.

When multiple people and organizations are working in a similar space, and they all try to take credit for something that any one of them might have accomplished alone, it results in every organization having inflated cost-effectiveness analyses, because the endline impact is subject to **double/triple counting**. Any CEA attempting

to model the value of meta-activities needs to take into account the direct impact and counterfactual replaceability of those actions in order to determine whether they are really tackling the limiting factors to direct work. Without accounting for these considerations, it will always *appear* more valuable to be a meta organization – but in the end, if everyone only did meta activities then no valuable work would get done.

Double counting is not only restricted to meta organizations, and can also happen when you are doing direct work. There will be many interacting actors whose work will make your intervention possible (e.g. you can't fortify without mills, or increase vaccinations without hospitals and vaccination camps) and in some cases the impact that you realize might have otherwise have happened anyway. Luckily, it's a little easier to avoid these sorts of errors when you can do a randomized controlled trial.

Assuming your impact lasts forever, or that no one else would do it

Generally, anything you build will eventually fall apart – either due to failure, or the world changing and moving on. Just because you got a farm to pledge to fortify their chicken feed now, doesn't mean they will keep fortifying it for decades. Just because you passed a government policy change, doesn't mean that the policy won't be reversed in the future.

On a similar note, it is often the case that someone else might eventually have implemented your intervention, had you not done so. So it's often best to model your impact as *speeding up the arrival* of an intervention.

Both of these considerations can be addressed using a time-discounting factor. For example, if you have hatched a plot that will save 100 children 5 years into the future, but with each passing year there is a 3% chance that some unforeseen circumstance will foil your plans, you should model your actions as saving $100 \cdot (1 - 3\%)^5 = 100 \cdot (97\%)^5 = 100 \cdot 86\% = 86$ children. In other words, since you want to apply a 3% yearly discount, you should multiply your projected impact by 97% for every year into the future.

Incorrect assumptions about trends and distributions

Not every distribution is a normal distribution. Many statistical techniques will go wrong if you assume something is a normal distribution when it isn't. Trends that seem linear will often hit diminishing returns eventually. Trends that seem exponential will often turn out to be sigmoid (S-shaped) curves.

Forgetting time costs, management costs, and operation costs

Staff time is expensive, not only in money but in opportunity costs. It's common for people to think that if something does not come with a large bill, it's free, but spending a lot of staff time on something is actually very expensive. It's also easy to focus on the costs of the intervention itself and forget that things like managing staff or registering a charity or processing donations and payments and reimbursements also cost money, in the form of occupying the time of paid staff members.

Overgeneralizing evidence across contexts, or using just one study

Findings from one country, or even one village, often **don't generalize** to other settings. On the other hand, a **single focused study** in a highly applicable and adjacent context is also shaky ground on which to establish an intervention. Unfortunately, evaluating evidence correctly is difficult, and experts won't necessarily agree on the right way to do it. The best way to avoid making this mistake is to learn a lot about your intervention and the particular mechanism by which it operates, so that you can make these judgments on a case by case basis at least as well as any other expert in the field. In the ideal scenario, you would be able to combine your own RCT with prior evidence from studies on similar interventions around the globe, but this evidence will not be available for all interventions.

CEAs that are too complicated, illegible, or don't get to the point

The more moving parts your CEA has, the more likely there is an error. It is extremely common for simple accounting and copy-editing errors to throw off the final estimate by an order of magnitude.

It is always possible to add more modifying factors to a CEA – and you shouldn't shy away from accounting for complexity! But in most CEAs, there are generally a small number of important figures and key assumptions that can *mostly* explain the final number.

You should be able to identify these major factors, and when explaining your CEA to other people you should be able to describe them in a few sentences and a simple equation. Unless you're dealing with something that genuinely *can't* be reduced to a simple explanation, don't let your stakeholders wade through a complicated spreadsheet or Guesstimate model to understand what's really going on, only to find that the important bits could have been explained in a few lines.

Resources

Guesstimate

Why we can't take expected value estimates literally (even when they're unbiased)
Equal Application of Rigor for EA Interventions

How meta-charity can dominate direct charity
EA risks falling into a "meta trap". But we can avoid it.
Triple counting impact in EA

The Generalizability Puzzle
Beware The Man Of One Study

4. Using Speculative CEAs to Make Decisions

Author:

Ishaan Guptasarma,
Curriculum Developer & Content Writer

In the course of running a charity, you will often be faced with decisions where you need to predict the cost-effectiveness of an action before you take it. Here are some practical examples of how one can use CEAs to speculate on how cost-effective an intervention might be, before actually gathering any data of your own.

Using quick CEAs to rule out ideas

Idea: Pigs become aggressive and injure each other when confined in monotonous environments, which necessitates further welfare-reducing practices such as tail docking (to preempt tails being bitten off). Adding **branched chain toys** to pig farms for enrichment purposes is cheap, and reduces these forms of aggression and other behavioral signs of distress.

Cost: Given the cost of maintaining the chains, we estimated at least \$0.50 per enriched pig-year.

Comparison: ACE estimates that The Human League averts 13–40 years of farmed animal life per dollar.

Moral trade-off: Would you rather spend a dollar to enrich a pig's cage for 2 years, or spend a dollar to avert 13–40 years spent in a factory farm for various animals?

Conclusion: Directly buying branched chain toys for pigs probably isn't as cost-effective as the alternatives.

This analysis has some problems! We don't know whether THL *really* averts 13–40 years of animal life per dollar. We don't know whether giving more money to THL will necessarily lead them to do any better than they are already doing. We might not have followed the same methodology as ACE and applied **equal rigor**, so comparing the numbers might be meaningless. We're committing the sin of **taking cost-effectiveness estimates literally**. We haven't carefully wrestled with moral questions (e.g. maybe you think 1 pig year is 50 times as important as 1 chicken year, and ACE mostly is averting chicken-years?). We don't know if \$0.50/enriched pig

year is the true cost.

However, **time spent analyzing things is a limited resource**. If we're just vetting initial ideas to decide what to rule out, this rough analysis might be enough information to come to the decision not to research this idea more thoroughly. This way we will have more time to consider more options, as well as more time to spend on interventions that actually look promising.

What if you decide an idea is worth taking a closer look at after all? Let's take a deeper look at some slightly more developed CEAs.

Using 2-hour CEAs to prioritize ideas for further research.

Cost-effectiveness analyses were the primary methodology used to select our [top 33 mental health interventions](#) this year. Ideas were ranked according to how well we expected they would improve people's answers on the [Satisfaction with Life Scale](#).

Let's compare two of the ideas that we considered for further research.

One of the ideas we considered was distributing anti-depressants. Here is the [2-hour CEA for distributing antidepressants](#).

One of the ideas we considered was lobbying for [positive education](#), which is the practice of incorporating mental self-care material into standard school curricula. Here is the [2-hour CEA for lobbying for positive education](#).

Comparing the costs

Here are the estimated costs of the two interventions we're comparing, side by side. You can follow along on the spreadsheets if you'd like to click the links. These images are just screenshots of the spreadsheet, lightly edited for clarity.

DISTRIBUTE ANTIDEPRESSANTS		POSITIVE EDUCATION POLICY	
Total intervention cost	150000	Total intervention cost	1550000
Fixed costs	150000	Fixed costs per year	310000
Co-founder salaries	50000	Co-founder salaries	50000
Operational costs	10000	Operational costs	10000
Staff	90000	Staff	90000
		<u>Lobbyists</u>	160000
Intervention costs	12		
Monthly cost	5.1	Years lobbying	5
<u>Cost per pill (India)</u>	<u>0.17</u>		
<u>Cost per pill (India)</u>	<u>0.17</u>		
Delivery/logistical cos	5		
Diagnostic costs	2		
Outcome / Scale			
Treated per year	25000		

Fixed costs. If a cost is fixed, then it means it will stay approximately the same even if the intervention scales up. This can include yearly salaries of key employees, plus some operational costs such as renting an office. In reality, the number of employees will often increase with the size of the organization, but proportionally speaking staff cost will often be more fixed. For example, staff salaries and the total number of staff generally do not increase if the number of beneficiaries increases from 20,000 to 25,000.

We've assumed that positive education policy will have large fixed costs because you would need to hire lobbyists. Lobbyists are typically more expensive and will want competitive salaries.

Intervention costs: This is the cost of the intervention itself – in this case, diagnosing the issue, buying the pills, and getting them to where we need to go.

In the case of distributing antidepressants, we made some guesses for how much it would cost to buy and get out each pill, and calculated the cost of one month of treatment. The pill costs are available online. If we had infinite time, we would have some better justifications for the delivery and diagnosis costs that we guessed – but this is only a 2-hour CEA!

Total cost – Distribute antidepressants: We've assumed that 25000 people get treated, so the **total cost** = **fixed cost** (150k/year) * **intervention cost** (\$12/month per person) * **treated per year** (25000 people) = **\$452.5k** over....

...Did you spot the error? We calculated the **intervention cost** at \$12/month, but we estimated 250,000 people *treated* per year, and then we multiplied those together. That means that if we intended our final “cost” to be “per year” then we are actually off by a factor of 12. This means that an important thing in the 2-hour CEAs we published was wrong!

When you make CEAs, especially quick ones, no matter how good you are, it is incredibly common to make errors such as these. Even GiveWell, who has had several different people review and recheck and re-recheck their spreadsheets for errors, once **allocated \$100k incorrectly due to a spreadsheet error**. This is one of many reasons it's important to publish your CEAs publicly and have others review them.

Luckily, we use many methods, not just CEAs, to keep our research process robust against these sources of errors. This is necessary because even if we never made any obvious mistakes (which of course we do), many things other than arithmetic errors can make CEA be orders of magnitudes off.

Let's keep going, and see what the consequences of this error are.

Total cost – Positive Education Policy: We've assumed that whatever benefits might occur, would occur within 5 years of lobbying. Therefore we simply multiply the fixed costs by five. **Fixed cost** (\$310k/year)***5 years** = **\$1.55m** over **5 years**.

Why are we looking at the cost over 5 years, rather than per year? You could argue that this artificially inflates the cost, or you could say that it makes more sense to talk about shifting the probability of policy change per year of operation. You could also argue that it's more appropriate to look at lobbying as a one time purchase of a process that takes 5 years and has a certain percent change of success, since lobbying doesn't “scale” and there's a specific policy objective to accomplish that does not necessarily require more funding once accomplished. It might be that similar policy initiatives tend to move on to a new project after 5 years of trying. There is not a clear answer. Reasonable people will disagree and there are all kinds of subjective judgment calls to be made when making a CEA!

Comparing the benefits

Here are the estimated benefits of the two interventions we are comparing, side by side. You can follow along on the spreadsheets if you'd like to click the links. These images are just screenshots of the spreadsheet, lightly edited for clarity.

DISTRIBUTE ANTIDEPRESSANTS		POSITIVE EDUCATION POLICY	
Total SWLS per year of treatment	54994	Total expected SWLS	5724326
		Average years before reform	5
Adherence rate	80%	Probability of bill success	35%
Treated per year	25000	Probability that bill proposed	25%
		Total SWLS per year	13235435.53
		Total pupils per year	8060000
SWLS per Treatment	2.750	State secondary school pupils	3330000
Study on paroxetine on SWLS	3.1	State primary school pupils	4730000
SWLS from treatment effect	2.39939	School leavers per year	644800
% in remission	45.40%	Average number of school years	12.5
Generalizability discount	50%		
SWLS average (Spain)	24.16	Final year carry over effect	1.41155896
SWLS for depression (Madrid)	13.59	SWLS per year of treatment pupil	1.5
		6-12 month after treatment effects	1.3
		0-6 month/active effect	1.5
		SD of SWLS	6
		Effect size 0-12 months	0.22
		Effect size 0-6 months	0.26

Using academic papers to estimate treatment effects: We used academic papers to estimate the change in subjective well-being for people receiving **antidepressants** and **positive education**, respectively.

Figuring out how many people are impacted: For antidepressants, we just made a guess about how many people we would be able to reach (25,000) and guessed that 80% of them would adhere to the program, based on dropout rates from similar things. For positive education, we estimated how many people would be in school in the UK, which is the region of the suggested policy change.

Figuring out chance of success: We assumed that our work would increase the chance of the bill being passed by 25%. (This is a notable, and perhaps arbitrary choice we made in this CEA. It's unfortunately really hard to guess how much more

likely one's work makes a policy change). We guessed that the bill would have a 35% chance of passing if proposed, based on **data from previous bills**.

Calculating impact

Impact = treatment effect * number of people impacted * chance of success

Antidepressant distribution

25,000 **treated/year** * 80% **adherence** * 2.750 **SWLS gains/treatment**
= **54,994 SWLS/year**

Positive education policy

25% **gain in bill proposal chance** * 35% **bill passing chance** * (8.06m **students** * 1.5 **SWLS gains per year** + 64,480 **school leavers** * 1.4 **after treatment effects**) * 5 **years**

Note the “5 years” of benefits presumed at the end. This presumably reflects an assumption that the policy change will last for 5 years before being repealed, or that the policy might have passed without our intervention 5 years later. Since they use the same cell in the spreadsheet, perhaps the underlying assumption is that the policy would hold for approximately as long as it took to lobby for it.

Converting to cost-effectiveness:

Cost-effectiveness is simply a division of benefits by costs. Distributing antidepressants averages **\$10.96 dollars/SWLS** (implying that each \$10.96 spent causes the sum of all beneficiaries' SWLS scores to be one point higher, were they to take the test) with **+\$4.4/additional person treated**.

But wait! There was an error in this CEA earlier, so we have to increase that price by a factor of 12. That would bring us to **+\$131.52 dollars/SWLS** and **\$52.8/additional person treated**. This is not as good, though still pretty good.

Meanwhile, positive education policy averages **\$0.27 dollar/SWLS**. Does that mean that positive education policy is much better than antidepressant distribution? It might be, but not necessarily! The assumptions that went into these CEAs are quite different, and it's hard to compare them side by side. Plus, both CEAs involved shaky assumptions: for example, it might be much too optimistic to think that you can shift the chance of a policy being proposed by 25%.

Don't rely only on CEAs!

If reading this sequence has left you feeling less confident in CEAs, it has accomplished its goal. The results of CEAs can feel very final and precise when all you are presented with is a simple number, but as you've seen, the process of making CEAs is actually full of judgment calls, and small mistakes and arbitrary choices can have outsized consequences.

Luckily, we don't use only CEAs to pick our interventions! There were ideas that did very well on CEA analyses that we did not recommend, and ideas that did relatively poorly that we did.



Image: Our *recommended ideas* had a wide range of CEA results

You can use many heuristics other than *cost-effectiveness analysis* to think about impact. Some of the ones that we use include *expert views*, *weighted factor models*, *informed consideration*, and *prior view*.

The purpose of making a speculative CEA is less to get “the right answer” and more to force yourself to think more rigorously about how much things are going to cost. Using a variety of methods, rather than only CEAs, ensures that our charity ideas

are robust to the many mistakes and random variations that occur when making a CEA.

Summary

Quick CEAs can be used to rule out ideas (but not rule them in!) by showing how some interventions are likely to be strictly worse than others. A quick analysis is often enough to inform you about whether it's worth looking further. Because this analysis must be done fast, it will have many mistakes and problems. However, this is acceptable because your time is limited when considering many options. You can do a more detailed CEA once you have a better idea of which options are likely to be worthy of further investigation.

For most CEAs, costs can be divided into “fixed costs”, which are the same every year, and “intervention costs”, which tend to scale as the organization grows. Calculating benefits tends to be a matter of the number of people you can impact and the effect of your actions on those people.

One huge issue with CEAs (even ones that you spend more time on) is the potential for basic mistakes, such as missing a decimal point or adding the wrong numbers. As your CEA develops and becomes more important, you should develop a simpler model (try to explain it in a few sentences or a short equation) which you understand from multiple perspectives and have done sanity checks on, so that you won't fall prey to arithmetic errors.

Resources

[Using a Spreadsheet to Make Good Decisions: Five Examples](#)

[Equal Application of Rigor](#)

[13 Common CEA Mistakes](#)

[What is an Impact Evaluation?](#)

[Why Location Matters When Picking an Intervention](#)

5. Using More Rigorous CEAs: New Incentives, Fortify Health, and CSH

Author:

Ishaan Guptasarma,
Curriculum Developer & Content Writer

Thus far, three organizations connected to Charity Entrepreneurship have received detailed third-party CEAs from GiveWell – [New Incentives](#), [Charity Science Health](#) and [Fortify Health](#).

New Incentives (NI) was the first organization to receive a GiveWell Incubation Grant. New Incentives uses cash transfers to incentivize mothers to vaccinate their children in the northwestern region of Nigeria. Patrick Stadler, co-founder of New Incentives, is the Director of Communications at Charity Entrepreneurship, and has been a mentor to many of our staff and incubatees.

New Incentives Sheet	Median	Andrew	Chelsea	Chris	Elie
calculation	results	subjective input	input to be updated by	more data	
Summary & Results					
New Incentives vs AMF	0.78	0.77	1.24	0.74	0.54
New Incentives vs GiveDirectly	5.15	4.30	14.12	10.89	4.90
Cost per equivalent outcome to an under-5 death averted	\$1,963	\$1,854	\$1,707	\$2,871	\$2,490
Percent benefit driven by preventing under-5 deaths		95%	98%	98%	97%
Percent benefit driven by consumption benefits		5%	2%	2%	3%
PCV (vaccine against pneumonia)					
Vaccination rate without the program: PCV 1		37%	33%	33%	37%
Vaccination rate without the program: PCV 2		30%	29%	29%	30%
Vaccination rate without the program: PCV 3		23%	24%	24%	23%
Vaccination rate without the program: exactly 0 doses of PCV		63%	67%	67%	63%
Vaccination rate without the program: exactly 1 dose of PCV		8%	4%	4%	8%
Vaccination rate without the program: exactly 2 doses of PCV		7%	4%	4%	7%
Vaccination rate without the program: exactly 3 doses of PCV		23%	24%	24%	23%
Vaccination rate with the program: PCV 1		52%	53%	43%	47%
Vaccination rate with the program: PCV 2		45%	49%	39%	40%
Vaccination rate with the program: PCV 3		38%	44%	34%	33%
Vaccination rate with the program: exactly 0 doses of PCV		48%	47%	57%	53%
Vaccination rate with the program: exactly 1 doses of PCV		8%	4%	4%	8%

Image Source: [GiveWell's CEA of New Incentives](#)

Charity Science Health (CSH) uses SMS text messages to remind people when it is time to vaccinate their children in several Indian states. The CE founding team conceived this charity using the same research process that generates our charity recommendations today. Its success was one of the factors that inspired us to start Charity Entrepreneurship.

Charity Science Health	Median	Andrew	Chelsea	Chris	Elie
Summary & Results					
Charity Science Health vs AMF	1.0	1.0	1.4	1.3	0.9
Charity Science Health vs GiveDirectly	7.1	5.4	16.0	19.0	8.1
Cost per equivalent outcome to an under-5 death averted	\$1,500	\$1,482	\$1,503	\$1,643	\$1,503
Percent benefit driven by preventing under-5 deaths		100%	100%	100%	100%
PCV (vaccine against pneumonia)					
Vaccination rate without the program: PCV 1		90%	90%	90%	90%
Vaccination rate without the program: PCV 2		89%	89%	89%	89%
Vaccination rate without the program: PCV 3		87%	87%	87%	87%
Vaccination rate without the program: exactly 0 doses of PCV		10%	10%	10%	10%
Vaccination rate without the program: exactly 1 dose of PCV		2%	2%	2%	2%
Vaccination rate without the program: exactly 2 doses of PCV		2%	2%	2%	2%
Vaccination rate without the program: exactly 3 doses of PCV		87%	87%	87%	87%
Vaccination rate with the program: PCV 1		94%	94%	94%	94%
Vaccination rate with the program: PCV 2		93%	93%	93%	93%
Vaccination rate with the program: PCV 3		92%	92%	92%	92%
Vaccination rate with the program: exactly 0 doses of PCV		6%	6%	6%	6%
Vaccination rate with the program: exactly 1 doses of PCV		1%	1%	1%	1%
Vaccination rate with the program: exactly 2 doses of PCV		1%	1%	1%	1%

Image source: *GiveWell's CEA of CharityScience Health*

Fortify Health (FH) works with mills in Maharashtra and West Bengal to fortify *atta* (a wheat flour commonly used in India) with iron, folic acid, and B12. We helped Nikita Patel and Brenden Eappan to found Fortify Health with a seed grant and mentorship. Brendan and Nikita's success showed us the power of mentoring talented people to start highly effective charities, and inspired us to launch Charity Entrepreneurship.

Results	
Units of value per dollar donated	0.04
Total units of value per \$10,000 donated	421.23
Outcome metric (select below)	
Cost per outcome as good as: averting the death of an individual under 5 — AMF	\$1,121
Contribution of each outcome to overall cost-effectiveness	
Percent of value coming from anemia morbidity averted	39%
Percent of value coming from cognitive benefits in children	8%
Percent of value coming from cognitive benefits in adults	53%

Image source: *GiveWell's CEA of Fortify Health*

This article will help you understand each of these GiveWell CEAs, with the goal of enabling you to use these processes to make key decisions about your charity.

How will understanding these CEAs help me start a charity?

In the previous article, we showed you how to **use speculative CEAs to make decisions**. However, when you are considering investing a lot of resources into a course of action, it is worth approaching the issue with deeper rigor.

Because these three organizations are closely linked to us, we have a lot of detailed information about the decisions they made leading up to becoming GiveWell-incubated organizations. We are, therefore, able to answer quite specific questions about how each organization works, and illustrate the connection between decisions that were made early on in the life of the organization to the organization's eventual cost-effectiveness as evaluated by a trustworthy and skeptical third party.

Understanding these CEAs lays the foundation to facilitate further conversations. For example, in our next article we'll cover how to **pick a location for your intervention**, with templates on how other charities chose their location. You will be able to see how making these choices prudently in the beginning led to high cost-effectiveness in later evaluations.

Comparing cost-effectiveness analyses for New Incentives and Charity Science Health

We'll start with an analysis of GiveWell's CEA of New Incentives and GiveWell's CEA of Charity Science Health, since both organizations are vaccination-driven and have very similar CEAs. This strong model equivalency enables an equal application of rigor when comparing the two interventions. All screenshots are from the New Incentives CEA, but the Charity Science Health CEA is similar. (Please note that the CEA model of New Incentives dates from 2017, and may not be the latest model GiveWell works with for the organization.)

Cost per equivalent outcome to under-five death averted (New Incentives/CSH)

	A	B	C	D	E	F	G	H	I	J	K	L
1	New Incentives Sheet	Median	Andrew	Chelsea	Chris	Elie	Isabel	James	Josh	Natalie	Sophie	Anyone
2	calculation	results	subjective input to be updated by more data						subjective input uncertain input			
3												
4	Summary & Results											
5												
6	New Incentives vs AMF	0.78	0.77	1.24	0.74	0.54	0.78	1.49	0.49	0.79	0.96	#VALUE!
7	New Incentives vs GiveDirectly	5.15	4.30	14.12	10.89	4.90	5.15	10.87	2.98	5.68	1.16	#VALUE!
8	Cost per equivalent outcome to an under-5 death averted	\$1,963	\$1,854	\$1,707	\$2,871	\$2,490	\$1,782	\$1,493	\$3,240	\$2,052	\$1,963	\$4,490
9												
10	Percent benefit driven by preventing under-5 deaths		95%	98%	98%	97%	96%	98%	93%	97%	80%	100%
11	Percent benefit driven by consumption benefits		5%	2%	2%	3%	4%	2%	7%	3%	20%	0%
12												

This is the most important value: the dollar amount this model suggests that New Incentives/CSH require to avert the death of one child under five years old, plus some additional benefit factored in as a result of the cash transfer (more on that later). As will soon become apparent, this value is not to be taken literally – its primary purpose is to compare similar organizations with a single metric.

Why are there multiple values? There are many subjective judgment calls to be made, and there are nine GiveWell staff members making these judgment calls. Anyone who wishes to make their own judgment call is invited to make a copy of the spreadsheet and can use the “Anyone” column.

(P.S. – from here on, hyperlinks will link to the specific cell or range being discussed when appropriate, so you don't have to go hunting for it in the spreadsheet.)

Modeling the cohort (New Incentives/CSH)

Model (without program)											
Visit 1 (BCG)											
Suppose a cohort of this size is eligible for the program just prior to going to		10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000
Deceased		0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Alive & immune to tuberculosis		1,897	1,897	1,707	1,897	1,802	1,897	1,818	1,707	1,897	1,897
Alive & susceptible to tuberculosis		8,103	8,103	8,293	8,103	8,198	8,103	8,182	8,293	8,103	8,103
Alive & carrying tuberculosis		0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Between Visit 1 (BCG) and Visit 2 (PENTA1, PCV1, (ROTA1))											
Number of weeks between Visit 1 and Visit 2		4	4	4	4	4	4	4	4	4	4
Event: some infants from group "alive & susceptible" contract tuberculosis											
Contracted tuberculosis in this period (new cases)		7.17	7.17	7.34	7.17	7.26	7.17	0.94	7.34	7.17	7.17
Still alive & susceptible		8,096	8,096	8,285	8,096	8,190	8,096	8,181	8,285	8,096	8,096
Event: some infants from group "new cases" die of tuberculosis											
New case & deceased in this time period		1.94	1.94	1.59	1.94	1.57	1.55	0.26	1.98	1.94	1.94
New case & survived (carrier)		5.24	5.24	5.76	5.24	5.69	5.63	0.69	5.36	5.24	5.24

This is a model of a hypothetical cohort of 10,000 people who might exist within an imaginary population targeted by this intervention. You may recall from your own medical history that different vaccinations must be administered at different ages, and some vaccinations require “booster shots”, requiring multiple trips to the clinic. GiveWell models this cohort as potentially visiting the clinic to be vaccinated 5 times between birth and age five. For each time period, they estimate the probability of various events, such as death and sickness from various causes, becoming immune to a disease, and so on. **These are modeled both with the program (treatment), and without the program (control), as a way of estimating the effect of the program.** (Once again, there are multiple values to the right because each GiveWell staff member is making their own judgment call.)

Parameters and assumptions (New Incentives/CSH)

In order to model the outcomes of the cohort above, we need to know some facts about the cohort demographics, including immunization status, mortality rate from various causes, and prevalence of various diseases. The parameters section specifies the databases providing the information regarding these variables. It also includes educated guesses about the predicted overall effect of the intervention on the cohort, based on data from RCTs of other similar interventions.

Parameters						
Use which incentive size?		Smaller	Smaller	Smaller	Smaller	Smaller
Use which vaccination rates?		North W	North W	North W	North W	North W
Prefer which vaccination rate data source?		SMART 2	DHS 201	DHS 201	SMART 2	DHS 201
Year of analysis		2018	2018	2018	2018	2018
Cap, vaccination rate without the program		95%	95%	95%	95%	95%
By default, assume what effect of the program (percentage point increase in		15%	20%	10%	10%	20%
By default, assume what effect of the program (percentage point increase in		15%	15%	10%	10%	20%

Pictured above: In the New Incentives CEA, databases covering particular regions of Nigeria are specified. These databases contain information about vaccination rates for multiple vaccine types in multiple regions. **The choice of database is important, as different databases provide different estimates** and some may be more accurate than others. In this case, the databases reflect immunization surveys conducted in Nigeria. The main criteria for prioritizing surveys are **recency** (as past surveys might be outdated) and **robustness** (which survey is likely to be more accurate, has a larger and more representative sample size, uses better methods, etc.).

The figures pulled from the regional databases will eventually be updated with more specific data. IDinsight is conducting an **RCT of New Incentives**, which at time of writing has collected data about **baseline characteristics of the population** and endline characteristics among the beneficiaries directly covered by the program. Using this data is better than consulting a general database, since general databases are estimates for entire regions, and do not necessarily consider the specific population that will visit clinics in which the intervention is being run.

Different staff members give different guesses regarding the degree to which the program will increase vaccination rates in the cohort. These estimates are taken from a reading of **Sato and Takasaki (2016)**, which found that another Nigerian cash transfer program had increased vaccination uptake of their cohort by 19 percentage points. Staff estimates are lower than 19 percentage points to discount for low generalizability. The value of this parameter determines how effective GiveWell staff believe the program will be at achieving its goal of improving vaccination rates.

The rough estimates given by staff regarding the effect of the program will eventually be replaced by IDinsight's RCT of New Incentives. At time of writing, data regarding the impact of New Incentives is unpublished – results are expected in summer 2020. Conducting this RCT is a **\$2m+ expense**, and the results will

determine whether New Incentives is cost-effective enough to be considered a GiveWell top charity.

Parameters						
Cost per enrollee		\$0.65	\$0.65	\$0.65	\$0.65	\$0.65
What proportion of the gap is reduced by SMS intervention		36%	36%	36%	36%	36%
External validity adjustment		100%	100%	100%	100%	100%
Replicability adjustment		100%	100%	100%	100%	100%
Total adjustment		100%	100%	100%	100%	100%
Adjusted proportion of the gap which is reduced by SMS intervention		36%	36%	36%	36%	36%

Pictured above: In Charity Science Health’s CEA, the parameters are a little different – instead of focusing on the raw percentage increase in vaccination, the focus is on the **reduction in the coverage gap**. For example, bringing immunization rates up from 70% to 85% would reflect a $(100\% - 85\%) / (100\% - 70\%) = 15\% / 30\% = 50\%$ reduction in coverage gap, since you’ve brought the number of *unvaccinated* people down from 30% to 15%. This parameter was used because it was thought to be a more ecologically valid way to apply to studies across contexts – you wouldn’t expect an intervention that “raises vaccination rates” by 30% in a low-income country to do the same in a high-income country, where the number of unvaccinated people will be lower to begin with. However, you might be able to generalize findings about closing a coverage *gap*.

Many of the assumptions for CSH parameters were taken from CSH’s [internal CEA](#). For example, the estimate for how much CSH could reduce the coverage gap was taken from the [Study results – Gap reduced](#) tab, which is a meta-analysis of RCTs undertaken for vaccination projects in other locations.

Baseline immunization and health status of beneficiaries ([New Incentives](#)/[CSH](#))

All-cause mortality		
Percentage point annual decrease in neonatal mortality		0.06%
Percentage point annual decrease in infant mortality		0.23%
Percentage point annual decrease in under-5 mortality		0.52%
Neonatal mortality (under 4 weeks; under 1 month)		3.76%
Infant mortality (under 52 weeks; under 12 months)		6.62%
Under-5 mortality (under 260 weeks; under 60 months)		13.30%
Weekly mortality risk, first month (28 days)		0.95%
Mortality between 1-12 months (period of 11 months)		2.75%
Weekly mortality risk, 1-12 months		0.06%
Mortality between 12-60 months		7.15%
Weekly mortality risk, 12-60 months		0.04%
Tuberculosis (disease)		
Tuberculosis: Chance of infection if susceptible, yearly		1.15%
Tuberculosis: Chance of infection if susceptible, weekly		0.02%
Tuberculosis: Case fatality rate		27%
Tuberculosis: Uncertainty adjustment for case fatality rate		100%
Tuberculosis: Adjusted case fatality rate		27%
Proportion of live neonates who are have been infected by and survived tut		0.00%
BCG (vaccine against tuberculosis)		
Vaccination rate without the program: BCG		32%
Vaccination rate with the program: BCG		47%
Vaccine efficacy in preventing disease		59%
Uncertainty adjustment: vaccine efficacy		100%
Vaccine efficacy, adjusted for uncertainty		59%

In order to use the databases and estimates from the “parameters” section to model the outcomes of the cohort above, we need to make various guesses about the practical implications of the assumptions and data. These considerations require a great deal of research, and there are many judgment calls to be made about which sources of evidence are trustworthy and which assumptions are viable. To give one of several examples, in this spreadsheet, **tuberculosis fatality rates are highly uncertain**, and different staff members have applied different “discounting” factors penalizing for uncertainty. As you begin to build spreadsheets of your own, you’ll notice that different **databases** often return different numbers because they are based on different studies.

As with the earlier section, this section features many areas where different staff members have selected different numbers. This is a good section to look at when considering what the best location or beneficiary population for this type of intervention might be.

You’ll notice that for the New Incentives spreadsheet, many of these cells are pulling from the databases specified in the “parameters” section.

Cohort deaths averted (New Incentives/CSH)

Total alive		8,956.66	8,974.31	8
Total dead		1,043.34	1,025.69	1
Under-5 Deaths Averted				
Under-5 deaths averted by the program		76.40	95.0	
Consumption Benefit				
<i>Without the program</i>				
Attend visit 1		3,215	3,215	
Attend visit 2		3,669	3,244	

This value is calculated simply by taking the cohort model that was discussed earlier, and subtracting the number of deaths both with (New Incentives/CSH) and without (New Incentives/CSH) the program.

Cohort consumption benefits (New Incentives)

<i>Consumption benefit</i>					
Baseline annual consumption per capita, USD		\$657.45	\$657.45	\$657.45	\$657.45
Additional costs incurred by each woman in the cohort		\$0.49	\$0.62	\$0.33	\$0.33
Cash transfers received by each woman in the cohort		\$5.64	\$6.15	\$5.17	\$4.99
Expected consumption benefit per woman in the cohort		0.008	0.008	0.007	0.007
Total consumption benefit for the cohort		78.00	83.71	73.35	70.64
Total consumption benefit for the cohort, in units of equivalent infant death		3.75	1.51	1.06	1.57
Total units of benefit per \$10,000 spent (personal units)		30.01	120.76	52.91	23.55

New Incentives gives people cash to incentivize them to vaccinate their children. GiveWell calculates the value of this cash to the recipients using the equation $\ln(\text{income} + \text{cash transfer}) - \ln(\text{income})$. This equation purports to measure the value of income, accounting for the fact that money does a person more good if they have less of it to start with. This is the “consumption benefit”. Since CSH isn’t giving out money, this only applies to New Incentives.

Moral weights (New Incentives/CSH)

13	Value Assignments				
14					
15	Value assigned to averting the death of an under-5 — Immunization		8.00	80.00	50.00
16	Value assigned to increasing $\ln(\text{consumption})$ by one unit for one person for		0.38	1.44	0.72
17	Ratio: averting the death of an under-5 is X times as good as increasing $\ln(c$		20.79	55.45	69.31
18					
19	Cost per outcome as good as: averting the death of an individual under 5 —		\$1,429	\$2,119	\$2,137
20	Cost per outcome as good as: averting the death of an individual under 5 —		\$7,977	\$24,091	\$31,261

GiveWell needs to compare organizations that have different types of benefits. These value assignments are called moral weights, and represent the trade-off between increasing consumption by “one unit” (giving people cash such that $\ln(\text{income} + \text{cash transfer}) - \ln(\text{income}) = 1$).

The median ratio, according to GiveWell staff ratings, is 20.79. If your value ratio = 20, that means you consider increasing people’s income by 20 “units” to be just as good as saving the life of one person under 5. To put this in terms of American currency and wages, a person who suggests a value ratio of 20 thinks that on some level of analysis (ignoring opportunity costs and so on), saving one baby brings as much raw value to the world as the quality of life increases that would result handing \$21k USD each to twenty people who live under the \$12k poverty line.

(This is because $x = \$21k$ in the equation $\ln(\$12k + x) - \ln(\$12k) = 1$, which matches our values ratio of 20 if we hand the money out to 20 people).

Costs and benefits per cohort (**New Incentives/CSH**)

Total Benefit				
Total benefit (equivalent under-5 deaths averted), per cohort		80.15	96.5	4
Costs				
Cost of cash transfers		\$56,382	\$61,478	\$51,
Administrative costs		\$84,878	\$93,878	\$75,
Costs to the government of additional vaccines and clinic load		\$7,352	\$9,323	\$4,
TOTAL COST		\$148,612	\$164,679	\$132,
Total units of benefit (units of equivalent under-5 deaths averted) per \$10,000 spent		43.15	468.79	174
Total units of benefit per \$10,000 spent (personal units)		345.168951	37503.3098	8708.24

This section describes the costs and the benefits for each 100,000 person cohort.

Benefits: The **cohort deaths averted** and **cohort consumption benefits** are combined using a single metric using **moral weights**, such that consumption benefits are expressed as morally “equivalent to” under-5 deaths averted.

Costs: The costs are estimated using real-world data regarding the charity’s operational costs, **holding fixed costs constant and scaling variable intervention costs** to the cohort size of 100,000 people.

Consumption & Cost Parameters						
Exchange rate, Naira per USD		300.00	300.00	300.00	300.00	300.00
Baseline annual consumption per capita		\$657.45	\$657.45	\$657.45	\$657.45	\$657.45
Cost of each clinic visit to the mother		\$0.67	\$0.67	\$0.67	\$0.67	\$0.67
Transfer per visit 1-4, Naira		500	500	500	500	500
Transfer for visit 5 (measles), Naira		2,000	2,000	2,000	2,000	2,000
Administrative cost per woman		\$18.00	\$18.00	\$18.00	\$18.00	\$18.00
Costs to the government, per additional individual vaccination visit		\$1.00	\$1.00	\$1.00	\$1.00	\$1.00

For New Incentives, as with most organizations, the main costs were administrative, at \$18 per mother enrolled in the program. The administrative costs here include both **fixed costs** (mostly administrative staff salaries), which remain constant as the program expands; and **scaling intervention costs** (mostly staff at each clinic), which grow as the program expands. This isn’t ideal, as we’d be interested primarily in scalable costs to calculate the value of our marginal dollar, but that makes it a conservative estimate. The second major cost for New Incentives was the 2,500 Naira (\$8.33 USD) cash incentive. The cash incentive does directly do good in the same sense that GiveDirectly produces good (and this benefit is incorporated into the model).

Parameters					
Cost per enrollee		\$0.65	\$0.65	\$0.65	\$0.65
What proportion of the gap is reduced by SMS intervention		36%	36%	36%	36%
External validity adjustment		100%	100%	100%	100%
Replicability adjustment		100%	100%	100%	100%
Total adjustment		100%	100%	100%	100%
Adjusted proportion of the gap which is reduced by SMS intervention		36%	36%	36%	36%

Estimates for the costs of Charity Science Health are sourced from **CSH’s internal CEA**. The \$0.65 per person number that GiveWell allocates for “cost per enrollee” can be found [here](#). The exact values are different, and perhaps CSH updated the analysis with additional data afterward.

The **combined salary of the administrative staff** is \$80–120k a year. This is a fixed cost, and in this case was actually estimated to decrease with scale as the intervention's infrastructure becomes more established and automated.

There are scalable intervention costs. These reflect the price of linking people's phone numbers to their vaccination schedules, so that appropriate reminders can be texted to them. CSH has considered several strategies to enroll people, and each of them comes with cost-effectiveness estimates. One way is to estimate the cost of hiring **local staff at hospitals** where vaccinations are administered and records are kept (\$3.1k for each staff, \$10k for staff managers). This method is estimated to acquire enrollees at a rate of \$0.25–\$0.35 per person.

These fixed costs and the scalable costs were placed into an equation and combined to estimate the cost per person. More details about how this was calculated can be found in the spreadsheet. A sanity check for this value was performed by looking at **the costs of another program** doing a similar thing. The service **vRemind** is a social enterprise that also provides vaccination reminders by text message in India, and their estimated cost is \$0.50–0.65 per person to acquire enrollees. This is similar to the values calculated by Charity Science Health, which helps create confidence that this CEA is realistic.

Cost-effectiveness (**New Incentives/CSH**)

Summary & Results				
New Incentives vs AMF	0.78	0.77	1.24	0
New Incentives vs GiveDirectly	5.15	4.30	14.12	10
Cost per equivalent outcome to an under-5 death averted	\$1,963	\$1,854	\$1,707	\$2,1
Percent benefit driven by preventing under-5 deaths		95%	98%	9
Percent benefit driven by consumption benefits		5%	2%	

We can now revisit this section of the spreadsheet, which we started with in the beginning, with a better idea of what is being said.

Cost per equivalent outcome to an under-5 death averted is calculated by dividing **benefits by costs**.

Percent benefit driven by preventing under-5 deaths/consumption benefits describes how much each aspect contributes to the total goodness of the intervention, for a given set of moral weights.

New incentives vs AMF and *New Incentives vs GiveDirectly* is a comparison of how cost-effective this intervention is relative to those organizations when put through a similar CEA, given a certain set of moral weights. These ratios will be one of several tools used to make **funding allocation decisions**.

Applying lessons: Fortify Health's CEA

GiveWell's CEA of Fortify Health is quite different from the other two. The intervention type is harder to model, and the CEA as a whole is less thorough and complete.

Before you read on, we encourage you to look over the CEA by yourself and see if you can identify parallels to some of the elements that we outlined above, such as various assorted benefits (e.g. deaths averted, consumption increased, and so on), moral weights, costs and benefits (with benefits consolidated into a single metric using moral weights), and cost-effectiveness relative to other organizations

Only some benefits are calculated (**Fortify Health**)

Items considered for inclusion	Included?	Rough best guess of effect size	Can it be objectively justified?	Ease of modelling	Consistency
		<15% = very low	0 = Not really	0 = It is impossible to model the effect with any reason	
	Color code	15-25% = medium	1 = Quantifying the effect is a gu	1 = It's unclear how we wo	1 = The effect would be equ
		25-50% = high	2 = Quantifying the effect requir	2 = The methodology is cle	2 = The effect is included in c
		>50% = very high	3 = There is strong direct eviden	3 = The methodology is cle	3 = The effect is included in
Iron fortification					
Anemia morbidity averted	Included	>50%	2	3	
Cognitive benefits in children	Included	25-50%	1	2	
Cognitive benefits in adults	Included	>50%	2	3	
Altered malaria risk	Excluded	<15%	3	3	
Benefits from averting neural tube defects due to folic acid fortification	Excluded	5%	3	3	
Possibility that Fortify Health's partner millers would have fortified anyway	Excluded	5%	3	2	
Possibility that the Indian government mandates fortification	Excluded	5%	1	1	
Maternal mortality reduction	Excluded	<15%	1	3	
Preterm birth risk	Excluded	<15%	1	2	
Increased birth weight	Excluded	<15%	3	2	
GI side effects	Excluded	<15%	3	3	

The three benefits that this CEA looks at are: 1) Anemia morbidity averted, 2) cognitive benefits in children, and 3) cognitive benefits in adults. The other benefits listed are not included in the analysis.

Why not? Aside from time constraints, recall from the [Equal Application of Rigor for EA Interventions](#) article that the more positive effects you account for, the more positive your intervention will look. There is always a very large number of potential positive and negative factors you could add to your CEA, and you don't want your result to be an artifact of what you happened to include. While Fortify Health adds iron, folic acid, and B12 to *atta*, GiveWell has only evaluated the benefits of the iron. Furthermore, only three of the benefits of iron are included here.

While not so explicitly stated, this was also true of the New Incentives and Charity Science Health CEAs above! Vaccination might have positive benefits for people who are over 5 years old, benefits relating to herd immunity, benefits to work capacity and therefore income, etc. But to keep comparisons equal between e.g. vaccination and malaria nets, these things are not included.

(PS – mousing over the little black triangles at the top right corners of cells indicate comments. This spreadsheet happens to contain many informative comments that can help you understand what's going on and where various numbers come from).

Intervention costs (Fortify Health)

Mill size (installed capacity of mill (MT/day))	2-10	10-20	20-60	60-100	100+
Annual costs	Source 1	Source 2			
Production costs					
Average installed capacity (MT/day)	6	15	40	80	125
Proportion of installed capacity actually utilised	0.69	0.69	0.69	0.69	0.69
Average actual production volume (MT/day)	4.1	10.3	27.5	55	85.9
% Atta that Fortify Health expects to eventually fortify in steady state	38%	57%	55%	61%	51%
MT/day of fortified atta per mill	1.5675	5.878125	15.125	33.77	43.828125
Total number of mills	2	6	5	4	3
Total MT/day of fortified atta	3.135	35.26875	75.625	135.08	131.484375
Total MT/year fortified atta (assuming 25 days/month operation)	940.5	10580.625	22687.5	40524	39445.3125
Multiplier to incorporate "overages" for less precise equipment	120%	115%	110%	105%	100%
Premix cost					
Premix cost per MT	\$1.80	\$1.73	\$1.65	\$1.58	\$1.50
Total premix cost	\$1,693	\$18,252	\$37,434	\$63,825	\$59,168
Equipment cost					
Average equipment cost for mill size	\$1,500	\$4,000	\$5,000	\$12,000	\$15,000
Total up-front equipment cost	\$3,000	\$24,000	\$25,000	\$48,000	\$45,000
Annual equipment depreciation rate	0.15	0.15	0.15	0.15	0.15
Total equipment cost per year	\$450	\$3,600	\$3,750	\$7,200	\$6,750
Total costs per annum (premix + equipment)	\$2,143	\$21,852	\$41,184	\$71,025	\$65,918

This spreadsheet carries information about intervention costs, and also serves as an index of the organization's activities – how many mills are fortifying, how much flour they are producing, etc.

For example, in column 1, if Fortify Health works at 2 mills with the capacity to process an average of 6 megatonnes (Mt) of *atta* per day, then under these

parameters it would be estimated to cost \$2143/year. The key assumptions in generating this number include: the capacity that the mills might actually run at (assumed 69%), the percentage of flour that actually gets fortified (assumed 38%), the cost of the pre-mix (a powder containing the nutritious compounds to be fortified – \$1.80/Mt, cheaper at larger scales), and an equipment depreciation rate (assumed 15%) to account for the fact that equipment doesn't last forever.

Fixed costs (Fortify Health)

Additional costs	Total
Staff costs	\$265,000
Mill operations	\$70,000
Business and publications	\$33,350
Travel	\$38,800
Registration, legal fees and fiscal sponsorship fees	\$37,311
Contingency	\$7,200
Monitoring and evaluation	\$100,000
Total additional costs	\$551,661

These are the fixed costs in relation to the ordinary operations of the charity itself. Currently, these comprise the bulk of the costs, but typically as an organization scales, these costs remain relatively constant and the intervention costs will begin to dominate. For most organizations, staff costs are a dominating factor. This overview also illustrates the investment in monitoring and evaluation, a substantial cost factor, but valuable to understand how effective a program is.

Calculating the number of beneficiaries (Fortify Health)

Beneficiaries reached	
Average consumption of atta per person per day (grams)	164
Average consumption of atta per person per year (grams)	59,999
Total quantity of fortified atta produced per year (metric tons)	114,178
Wastage adjustment	22%
Total quantity of fortified atta produced per year after wastage (metric tons)	89287
Total beneficiaries	1,488,151
Annual cost per person (nominal USD)	\$0.51

Unlike the previous CEAs, Fortify Health does not use cohort modeling. Instead, the total number of beneficiaries is calculated by figuring out how much flour the fortified mills produce and estimating how many people it would plausibly take to eat that amount of flour.

This is because, in the case of vaccination, we're largely interested in seeing if we should apply more funding to increase the number of people vaccinated through the program. The actual number of people being vaccinated is something to measure, not something to estimate. On the other hand, for fortification-oriented activities, we're primarily interested to understand if we should apply more funding to expand the number of mills that are enrolled in the program. The number of people reached and the amounts of the *atta* that they eat is a crucial consideration and can't easily be measured directly, only estimated.

If we can estimate the degree to which the average consumer is iron-deficient, and the benefits of them consuming this fortified *atta*, then we can estimate the cost per person of alleviating cases of anemia via fortification. Under the estimates and assumptions used in this CEA, this amounts to \$0.51 per person. We can also use this data to make decisions about how to reach even more people, and reduce this price.

Baseline iron deficiency among beneficiaries (Fortify Health)

	Children with any anemia	Women with any anemia	Men with any anemia	Average adult	Total population	Projected number of mills
Maharashtra	53.8	48	17.6	32.3	112374333	15
West Bengal	54.2	62.5	30.3	45.9	91276115	5
India	58.5	53.1	23.3	37.8		

In order to understand the effects of averting iron deficiency, we need information about how iron-deficient the beneficiaries are in the first place. This section uses state-level data and extrapolates existing trends to estimate the rate of anemia in the beneficiary population without the intervention, as well as its projections going forward in future years if the intervention ceases.

Cohort morbidity averted (**Fortify Health**)

Population-wide iron fortification	Active
Anemia morbidity averted	
Annual YLDs per person attributed to dietary iron deficiency in India	0.0054
Relative risk of iron deficiency with iron fortification from academic literature	48%
Implied discount factor for difference in effectiveness of FH's program compared to academic literature	115%
Adjusted relative risk of iron deficiency with iron fortification in FH's program	40.0%
Internal validity adjustment for iron deficiency	90%
External validity adjustment for iron deficiency	100%
Anemia YLDs per person 0-14 years averted by iron fortification	0.0029
Value assigned to 1 YLD averted	2.8
Units of value per beneficiary	0.0084

The impacts of anemia and the benefits of averting it are quantified in terms of **Years Lost to Disability (YLD)**. Moral and epistemic weights are applied to the value of those benefits so that they can be compared to averting under-5 deaths and increasing income.

Given the baseline rate of anemia, people who eat flour fortified by Fortify Health are 40% as likely to get anemia as people who are completely unfortified. In other words, they are $1 - 40\% = 60\%$ less likely to get anemia, although this number is multiplied by 90% to account for publication bias.

For each person spared of anemia, 0.0054 YLD are averted. This means that in total, $90\% * (1 - 40\%) * 0.0054 = 0.0029$ YLD averted, for every \$0.51 we spend. Each YLD is assigned 2.8 “units of value”. If you compare the **moral weights**, that’s 1/30th of the 85 “units of value” assigned to averting the death of a child under 5.

Cohort cognitive benefits (**Fortify Health**)

Cognitive benefits in adults	
Proportion of cohort of income-earning age	64%
Proportion of working-age adults that have dietary iron deficiency	23.8%
Increase in IQ from iron supplementation in adults with anemia (standard deviations)	0.35
Internal validity adjustment for IQ (adults)	65%
External validity adjustment for IQ (adults)	100%
Adjustment for lower effectiveness of iron fortification vs. supplementation	75.9%
IQ points per standard deviation	15
Increase in IQ from iron fortification in average adult (points)	0.39
Percentage increase in wages/consumption for every 1 point increase in IQ	1.0%
Multiplier for resource sharing	2
Percentage increase in consumption from iron fortification in average adult	0.79%
Increase in annual $\ln(\text{consumption})$	0.0078
Value assigned to increasing $\ln(\text{consumption})$ by one unit for one person for one year	1.4
Units of value per beneficiary	0.011

The impact of anemia on IQ, and of IQ on consumption, are quantified so that averting anemia can be compared to cash transfers. Pictured here is the analysis performed for adults. A similar analysis was done for the gains experienced by children.

The benefits of increasing IQ are attributed entirely to the effects that IQ has on consumption. In this system, an intervention which increases someone's IQ such that they would be able to earn more would be considered just as good as simply giving someone the value of those additional earnings as a cash transfer.

As you keep reading, you'll notice that this means iron fortification is estimated to generate several times more income for adult beneficiaries via IQ gains alone than they would by simply being given the amount of cash that goes into fortification. As long as certain practical assumptions hold, this means that fortification will be more cost-effective compared to cash transfers *regardless* of which moral weights are inputted into the spreadsheet.

Moral weights (**Fortify Health**)

Moral Weights	Active
Consumption Increases	
Value assigned to doubling consumption for one person for one year	1.0
AMF	
Value assigned to averting the death of an individual under 5 — AMF	47
Value assigned to averting the death of an individual 5 or older	85
Malaria Consortium	
Value assigned to averting the death of a 3- to 59-month old child — SMC	48
Helen Keller International (HKI)	
Value assigned to averting the death of a 6- to 59-month old child — HKI	49
YLDs	
Value assigned to 1 YLD averted	2.8
Derived value assignments	
Value assigned to increasing $\ln(\text{consumption})$ by one unit for one person for one year	1.4

These are the moral weights assigned to averting Years Lost to Disability, averting death at various ages, and increasing consumption (income). Moral weight to increasing IQ was not assigned, because gains in IQ were converted to gains in income.

Cost-effectiveness (Fortify Health in general and vs. GiveDirectly)

Intervention cost	
Annual cost per person (nominal USD)	\$0.51
Results	
Units of value per dollar donated	0.04
Total units of value per \$10,000 donated	421.23
Outcome metric (select below)	
Cost per outcome as good as: averting the death of an individual under 5 — AMF	\$1,121
Contribution of each outcome to overall cost-effectiveness	
Percent of value coming from anemia morbidity averted	39%
Percent of value coming from cognitive benefits in children	8%
Percent of value coming from cognitive benefits in adults	53%

Total units of value per \$10,000 donated with iron fortification	421.2
Total units of value per \$10,000 donated with GiveDirectly (according to the linked CEA)	38
Iron fortification vs. Cash	11.0

The various categories of benefit are converted into a single metric using moral weights, and the cost-effectiveness of these is compared to other organizations.

As with the CEA for New Incentives, these benefits are segregated by outcome – so if you want to directly compare giving cash to adults vs. increasing income via iron-induced cognitive benefits to adults, you could say that $11 \times 53\% = 5.83$. This roughly means that fortifying iron does an estimated 5.83 times more good than giving cash to beneficiaries directly, even when only considering the degree to which an adult might earn more cash due to not having anemia-induced cognitive difficulties. (Overall, iron fortification is 11x cash transfer if you include the value coming from anemia morbidity averted in children and cognitive benefits in children. This sum is highlighted in the bottom column).

Summary

The CEAs for all three organizations have certain features in common. Choices are made about which outcomes to model. Fixed costs and intervention costs are considered, generally separately. Statistics about the base rate of disease in the region of the intervention play a large role in estimating the effect of the intervention. Models are used to combine base rate statistics with known aspects of

the intervention to estimate the positive benefits. Finally, moral weights are used to compare different benefits to each other and to consider multiple benefits as a single unit.

Resources

[New Incentives CEA Winter 2017](#) by GiveWell

[\[Early 2017\] Charity Science Health CEA](#) by GiveWell

[Steady-state Fortify Health iron fortification CEA \[June 2019\]](#) by GiveWell

[Shared copy CEA 2017](#) internal CEA by CSH

PS – if there is any GiveWell evaluated organization that you’d like to see this and similar data on, GiveWell sometimes has useful “all content on X” pages. For example, if you wanted to know about CSH, you could do a google search for “site: GiveWell.org all content on Charity Science Health.” You will find links to CEAs, grants, conversation notes, and other useful information. If you can’t find such a page, it often helps to look at indexes of grants (e.g. [discretionary](#) and [incubation](#) grants) since most grants include a CEA as part of the justification for the grantmaking process.

[Equal Application of Rigor for EA Interventions](#)

[Approaches to Moral Weights: How GiveWell Compares to Other Actors](#)

[Why we can't take expected value estimates literally \(even when they're unbiased\)](#)

Additional help: [GiveWell’s cost-effectiveness analyses](#) are posted every year. To understand them, I recommend watching [this explanatory video](#) up to the 16:49 mark and following along on [this document](#) of the 2016 top charity cost-effectiveness analysis. The reason to use the 2016 version for learning your way around is that not all documents from the later version have been released. A more recent written explanation can be found [here](#) (but note that some of the links are to internal documents that you will not have access to).

Reference:

[Sato and Takasaki \(2016\), Peer Effects on Vaccination: Experimental Evidence from Rural Nigeria](#)

6. How CE Uses Cost-effectiveness Analysis

Author:

Joey Savoie,
Co-founder & Director of Strategy

Charity Entrepreneurship's research process involves **multiple stages** and four methodologies: **informed consideration** (IC), **expert views** (EpV), **weighted factor models** (WFM), and **cost-effectiveness analysis** (CEA). In this article, we explain why and how Charity Entrepreneurship uses CEAs as part of our research process, and discuss its strengths and weaknesses as a methodology.

A CEA consists of one or more calculations estimating the ratio of the cost of a given intervention relative to its impact. CEAs are particularly useful because they allow us to quantitatively compare different interventions. However, CEAs can be prone to errors and can fail to adjust for **prior views**.

Charity Entrepreneurship uses CEAs during three stages of our research at different levels of depth: first five minutes, then two hours, and finally twenty hours.

- First, we conduct an **idea sort** to sift out the top thirty or so ideas from our initial list of hundreds. We assess each idea using all four methodologies (i.e., CEA, EpV, WFM, and IC), spending twenty minutes per idea.
- Next we produce a **prioritization report**, spending two hours researching each idea from the idea sort and ordering them according to their promise. We use a different methodology for each cause area: CEA is used for our **mental health** ideas. At this stage we only apply one methodology to ensure sufficient depth.
- Finally, our eighty-hour assessments of the top interventions involve all four methodologies.

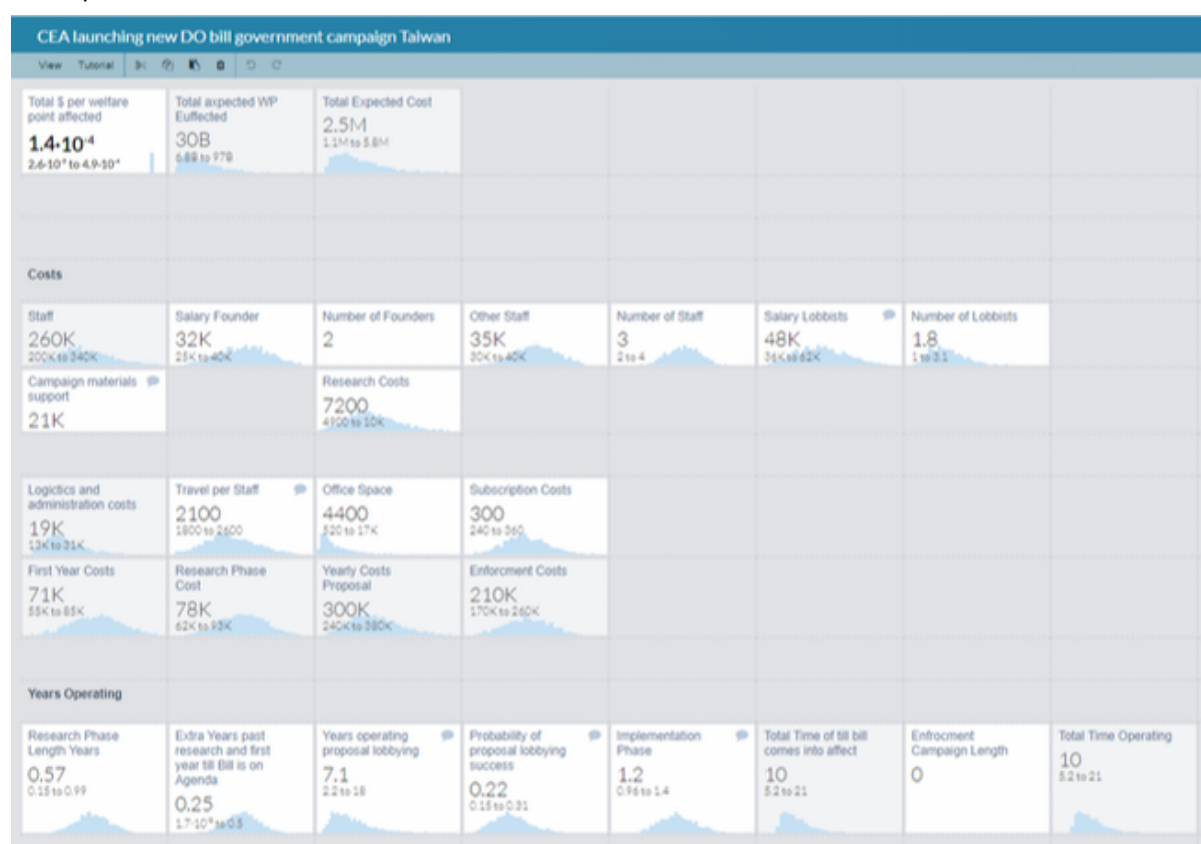
To conduct a CEA, lead researchers compile different key parameters and factors affecting the impact of an intervention into a sheet, and use these to calculate the ratio of cost to good done. Depending on the length of time allocated, the model will involve more or fewer parameters and factors, and will use more or fewer external sources to estimate them.

What is a cost-effectiveness analysis?

A cost-effectiveness analysis is a form of analysis most commonly used in economics, health economics, and charity evaluation. It consists of one or more calculations and results in a **ratio of the cost of a given action or intervention relative to its impact**. Costs are generally measured in dollars, with impact often measured in something like **DALYs** or lives saved. More cost-effective interventions generally have a lower ratio of costs to good done, and are considered better than less cost-effective interventions, all else being equal.

Summary: Cost-effectiveness is a measure of how much cost a given activity or action entails compared to the amount of good done by that activity or action. (1)

Example CEA:



Cost-effectiveness models vs. true intervention cost-effectiveness

It is important to distinguish between the **true cost-effectiveness** of an action and the **modeled cost-effectiveness**. The true cost-effectiveness of an action – if known

– would be a highly relevant metric and could be weighted very heavily when making a decision. However, the closest an evaluation-focused organization can usually get to ascertaining the true cost-effectiveness of an intervention is through constructing a model, which is almost by definition an imperfect estimation. This is because we often lack important data about the world, or a sufficient amount of it. Models can certainly be helpful and can be used as a type of evidence that an intervention should or shouldn't be considered. However, we think that it's important to take into account data from other types of models as well, given the shortcomings of cost-effectiveness models.

Why are cost-effectiveness analyses useful?

When attempting to compare the effectiveness of different interventions, it can be useful to create a formal and detailed model with a single endline number, i.e. one unique final result.

Benefits of CEAs (listed here from strongest to weakest):

- **Clearly connect to endline goals**
- **Can be used to compare interventions that are otherwise difficult to compare**
- **Allow formal sensitivity analysis**
- **Encourage quantitative analysis more broadly**
- **Are underutilized in many evaluations**
- **Give a transparent picture of the evaluator's rationale**
- **Are a respected tool in multiple fields**
- **Consider scope**
- **Reduce some biases**
- **Can lead to novel conclusions**

Clearly connect to endline goals: Ultimately, doing the most good per investment is our goal. A CEA may be an imperfect model, but it speaks directly to our key question. Compared to other models such as consulting experts or a weighted factor model, it has the clearest theoretical correlation with good done, even if in practice there are model errors that weaken it.

Can be used to compare difficult-to-compare interventions: Doing an analysis that results in a ratio is useful because it allows for a direct numerical comparison to be made. In other words, CEAs provide a way to quantitatively compare

interventions that may seem qualitatively incomparable (e.g. from different cause areas).

Allow formal sensitivity analysis: A sensitivity analysis can locate the most important assumptions, variables, and considerations affecting the endline conclusion. A formal sensitivity analysis can be done quickly and easily on a CEA, showing the key parameters that are the most important to get right. In other words, it identifies the factors that have the most substantial effect on the impact.

Encourage quantitative analysis more broadly: By default, most people (including experts) do not think in quantitative terms. For example, when asked if an event will happen, most people think of this as a binary question (yes/no) rather than thinking about the probability of the event happening. CEAs require quantitative inputs for each variable, which encourages quantitative thinking and calibration (e.g. an event being 20% vs. 80% likely).

Underutilized in many evaluations: Likely partially because both quantified consideration in general and CEAs take both considerable amounts of time and require a decent mathematical understanding, CEAs are often unused in situations that could benefit from them. Often in charitable areas with less established research bases, individuals trained in formal quantitative methods will have a very low CEA utilization. This allows CEAs to add a useful and unconsidered or under-considered viewpoint.

Give a transparent picture of the evaluator's rationale: Cost-effectiveness models provide a high level of transparency of thought. Since each input is identified and clearly quantified, an outsider can quickly see where assumptions are being made and can therefore more easily assess the validity of the conclusions.

A respected tool in multiple fields: Experts in many fields are in strong agreement that CEAs are a useful tool. These include experts in economics, medicine, and – most relevant for our purposes – [charity evaluation](#).

Consider scope: A major concern with many models is that scope is frequently not taken into account. If one charity has the potential to grow one thousand times bigger than another charity, a different type of model may not successfully reflect that it could be one thousand times more important to start the former charity over the latter. Humans are notoriously bad at properly understanding scope. (2)

Reduce some biases: CEAs are less susceptible to certain human biases that affect other analyses. For example, a well-used CEA can reduce the **base rate fallacy**, **conjunction fallacy**, and **hyperbolic discounting**.

Can lead to novel conclusions: CEAs can often lead to unintuitive conclusions and can, thus, lead to the consideration of new ideas or approaches that might have been quickly ruled out by other methodologies or “commonsense” approaches.

Why doesn't Charity Entrepreneurship rely exclusively on cost-effectiveness analyses?

Concerns with reliance on CEAs in charity evaluation have been discussed in depth in other posts, with the most comprehensive coverage of the theoretical concerns outlined by **GiveWell** and the most comprehensive coverage of the practical concerns by **Saulius Šimčikas**.

In their post “**Why we can't take expected value estimates literally (even when they're unbiased)**”, GiveWell states:

“The mistake (we believe) is estimating the ‘expected value’ of a donation (or other action) based solely on a fully explicit, quantified formula, many of whose inputs are guesses or very rough estimates. We believe that any estimate along these lines needs to be adjusted using a ‘Bayesian prior’; that this adjustment can rarely be made (reasonably) using an explicit, formal calculation; and that most attempts to do the latter, even when they seem to be making very conservative downward adjustments to the expected value of an opportunity, are not making nearly large enough downward adjustments to be consistent with the proper Bayesian approach.”

Flaws of CEAs listed here from strongest to weakest:

- **Subject to the “optimizer's curse”**
- **Necessarily involve value judgments**
- **Model uncertainty**
- **Prone to mistakes**
- **May not be generalizable to other contexts**
- **Make it hard to model flow-through effects**
- **Can be misleading in many ways**
- **The interventions we analyze are somewhat preselected for cost-effectiveness**
- **Subject to researcher bias**

- **May bias you towards interventions with more measurable results**
- **90% confidence intervals can be misleading**

Subject to the “optimizer’s curse”: All estimates are prone to error, and these errors compound. An intervention whose CEA yields a high cost-effectiveness is more likely to have had errors in its favor. This means that the most and least cost-effective interventions are likely to regress to the mean upon further examination. Overweighting CEAs in our decision-making could lead us to neglect good opportunities that did not have as many favorable errors. This is less of a problem in richer information environments.

Necessarily involve value judgments: It is surprising how much value judgments can differ. For example, GiveWell **assumes** that the "value of averting the death of an individual under 5 [years of age]" is 50 times larger than the value of "doubling consumption for one person for one year." Reasonable estimates could be as large as six times this number, using life-satisfaction years. If all value judgments are subjective preferences that vary among individuals, then CEAs are only generalizable insofar as the researcher’s values align with the reader’s.

Model uncertainty: Cost-effectiveness models are necessarily simplifications of reality. This is both a strength and a weakness. Although it allows us to get a clearer understanding faster, it also means that they do not accurately capture reality. Adjustments in the variables used will change the final value of the CEA. One way to combat this is to create several models and see if they converge.

Prone to mistakes: Mistakes are inevitable, due to human error and/or poor information quality. Although small mistakes usually only translate to small problems on their own, these mistakes compound in a multivariate model, thus exaggerating the consequences. For example, GiveWell once found five separate errors in a DCP2 DALY figure for deworming that contributed to an overestimation of the intervention’s cost-effectiveness by one hundred times.

May not be generalizable to other contexts: Some CEAs rely heavily on randomized controlled trials (RCTs) for their data, and in some cases, this can be problematic. If an RCT was conducted in one particular region or with one particular method, the effect size may change dramatically in different regions or with other methods.

Make it hard to model flow-through effects: Researchers have written that it is difficult to properly model flow-through effects in CEAs. Indeed, a common tactic is

to ignore flow-through effects entirely. There are solutions to this problem; however, they all take vast amounts of time or are prone to error.

Can be misleading in many ways: If researchers fail to consider important factors or are not transparent in their reasoning, CEAs can yield misleading results. For example, if a CEA concerns an expected value, the probability of success must be clarified. If only pure expected value is reported, there is no difference between a 50% chance of saving 10 children and a 100% chance of saving five children. This would fail to consider any level of risk aversion.

Interventions we analyze are somewhat preselected for cost-effectiveness: As the sources for our charity ideas were largely from within the EA community, these ideas will have been created with cost-effectiveness in mind. For our longer reports, we will have already narrowed down the ideas to more promising interventions, and thus, the variance will be lower. This means that random error will account for more of the variance, making CEAs a weaker tool.

Subject to researcher bias: CEAs are resistant to certain biases, but are susceptible to others. If the researcher conducting a particular CEA has a favorable view of the intervention, for example, he or she may (consciously or unconsciously) bias the results in its favor. A researcher's desire to find novel, cost-effective interventions may also have this result.

May bias you towards interventions with more measurable results: Effects that are difficult to measure may increase the error rate or be neglected. This can lead to an underestimation of the effectiveness of interventions with hard-to-measure outcomes.

Ninety-percent confidence intervals can be misleading: Depending on how well calibrated researchers are, the worst-case scenario, the best case and 90% confidence interval (CI) may be incorrect. CIs are particularly susceptible to this, as we are likely to underestimate the range of uncertainty that is actually accurate. Worst case and best case are no better, as these may rely on many unlikely events happening, meaning the probability of either occurring is minimal.

How much weight do we give our cost-effectiveness analyses?

Given that CEAs have many benefits and flaws, it is important to use them only in conjunction with other methodologies. CEAs are one of the four components of our evaluation process. We also consider the convergence of these components, i.e. in

what direction the different models point overall. We expect to weight our CEAs more than 20% but likely not more than 33% in our overall assessments (depending on the charity/intervention). We expect our CEAs to be more useful in areas where quantitative differences can be very large and where analysis based on our other evaluative criteria is less reliable.

How does charity entrepreneurship use cost-effectiveness analyses?

Software

We considered several software programs and combinations for our cost-effectiveness modeling. The two easiest to rule out were Google Documents for back-of-the-envelope calculation (BOTEC), and STATA, a complex modeling software. These had either too little or too much complexity for our purposes. We considered Google Sheets and Guesstimate in greater detail, and discuss those below.

Google Sheets: Google Sheets is fast and simple to work with, and it is easy to understand without much prior or complex knowledge. While spreadsheets are a common way to generate number-based models, they lack a few of the key features we need for our CEAs.

Guesstimate: Guesstimate is a less commonly used system, but has advanced Monte Carlo and sensitivity analysis features. It is too slow to use for very quick CEAs, but can be handy for models with high levels of uncertainty.

How we use the two in combination: For our 20-hour CEAs, we create them in a spreadsheet and then remodel the data using Guesstimate for both sensitivity analysis and simulated endline point estimates. The use of two models decreases the odds that an error in one model will have a very significant effect on the overall outcome, particularly since the software packages require somewhat different formatting.

Formatting

We use consistent formatting across all of our CEAs, and have tried to keep it somewhat consistent with GiveWell's formatting. This way, anyone familiar with GiveWell's CEAs will have an easier time understanding ours (and vice versa).

Color coding

Certain cells are color-coded to reflect the sources of those numbers.

Yellow: Value and ethical judgments

These numbers could change if the reader has different values from the researcher. For example, reasonable people could disagree about the answer to the question “How many years of happiness is losing the life of one child under five worth?”. When making these judgments, we generally consult the available literature, but there often is no clear, consistent, agreed-upon answer.

Green: Citation-based numbers

These numbers are based on a specific citation. If we found and considered multiple citations, the best will be hyperlinked to the number, and the others will be included in the reference section. If a number is an average of two other numbers, both numbers will be entered into the sheet, and the average will become a calculated number with a different color format.

Blue: Calculated number

These numbers are calculations generated from other numbers within the sheet. Calculated numbers involve no more than five variables for readability as well as easier sanity checking. Generally, it is harder to make errors if a higher number of subtotals are created instead of one very large, multi-variable calculation at the end.

Orange: Estimated numbers

Sometimes, no specific numbers can be found for a parameter. In this case, the number is estimated by one or more staff members. These estimates will often be the numbers within a CEA that we have the lowest confidence in.

Discounting

Discounting is a term we use for estimates that are affected or discounted by a factor not considered in the direct number. We try to keep our discounting clear and separate from the original number in the CEA, as these discounts are generally subjective. Discounting is common and can be seen in many other detailed CEAs (for example, GiveWell’s). The items listed below are not the only types of discounting used in our models, but they are some of the most common ones.

Evidence discounting: If a source of evidence suggests one number but the source is extremely weak, we might apply a certainty discount to it. This is based on the assumption that, in general, numbers regress as they get more certain. Thus, using a very weakly evidenced number in one estimate and a strongly evidenced number in another will systematically favor the areas with weaker evidence, as these numbers will be more positive.

Generalizability discounting: Often, sources will be based on a situation that is not identical or even similar to the situation we are considering when using a source. For example, if a study was run in one country, the results will not be identical if it were run in another country, even if all other factors are held constant. Thus, when generalizing evidence more than is common in our other comparable CEAs, we apply a generalizability discount.

Bias discounting: If a citation comes from a source that we suspect has some sort of bias, we might discount this number. For example, every charity has a strong incentive to make their program and progress look better. Thus, charity-reported numbers tend to be far more optimistic than the same activity analyzed by a study or outside actor.

Time discounting: Time discounting is the practice of discounting future benefits compared to immediate effects. Even with zero time preference, in terms of utility, it can still make sense to discount based on time. For example, income in the near term can be invested and used for increased consumption in the future. Additionally, there is always some probability that an accidental death will occur before the future utility is realized, and therefore, it is worth less in the present.

Organization

Each charity idea will have its own CEA sheet named after the charity idea being evaluated within the sheet.

Summary sheet: The first sheet will be a summary sheet that will allow quick comparison between the charities and will describe the three factors that could most change the CEA, as determined by a sensitivity analysis, and the factors considered the least certain by the CEA's creator. The summary sheet will include two endlines. One is a metric that is easily understandable and directly connected to the intervention – for example, “number of chickens’ life years lost from being in a caged vs. a cage-free system.” The other endline will be a cross-comparable metric that can be used across the entire cause area. This metric can be used to determine

which interventions look most cost-effective in a given area. There is also a column to describe the overall uncertainty level, which is the CEA creator's estimate of how confident we are in this CEA relative to others within the cause area.

Row, columns, and sections: The spreadsheet will generally be read across a given row with the first non-empty column containing a title or description. The column will generally be consistent across multiple rows. Specific sections will be put into boxes to increase readability. Generally, the last column and row of a given section will be used for notes or description. Every CEA will generally have a benefits, costs, and **counterfactuals** section.

Optimistic, pessimistic, and best guess: Throughout the spreadsheet, an optimistic, pessimistic, and best-guess estimate will be identified. The most time will be put into the best-guess numbers. The endline summary will be generated using a Monte Carlo simulation. The optimistic and pessimistic estimates will be used for the range of the 90% confidence interval. The relative position of the best-guess within this range will be used to determine the curve.

Sensitivity analysis: A sensitivity analysis will be conducted on each CEA to determine which factors most affect the estimate. The CEA creator then pulls out the factors that both have a large effect and seem more likely to change based on new information or research. These factors will be listed on the summary sheet.

Referencing: The most important or relevant reference will be linked to the cell. All references will be stored in a section of the reference sheet. Each cause-level CEA spreadsheet will have a consistent reference page as its last sheet. Our tracking of references will be done consistent with our system that is used across other methodologies.

How long are charity entrepreneurship's cost-effectiveness analyses?

Overview

At Charity Entrepreneurship we use four different CEA lengths throughout our charity idea evaluation process. CEAs are done in multiple lengths for two main reasons. The first is that there are simply too many ideas to conduct deep CEAs on each one. Our initial brainstorming often results in hundreds of ideas, so conducting

even a 10-hour CEA on each would require multiple years of research time. The second reason is that we want to see how the results of shorter CEAs compare to those of larger, more intensive CEAs. If a two-hour CEA yields results that are inconsistent with our endline recommended charity ideas, this could suggest that the CEA methodology is not effective for narrowing down ideas.

Five-minute CEAs

The goal of a five-minute CEA is to very quickly get a quantitative sense of the cost-effectiveness of an intervention. This CEA only considers key factors and only uses estimates based on intuition or citations that can be found in a very quick internet search.

At this level of CEA, results can vary by orders of magnitude. Therefore, high scores are considered to be almost certainly caused by errors. One way we account for this is by evaluating interventions using percentiles rather than raw CEA score. This limits the effect that errors can have on the overall score an idea receives and prevents the Z-score distribution from being too skewed by a particularly high or low CEA. One drawback of making these assumptions and corrections is that this may cause ideas that do have remarkably high cost-effectiveness to be undervalued. However, we believe that such ideas should also score well on other metrics.

We are skeptical about whether doing such brief CEAs yields helpful results. However, we use them in conjunction with other approaches, and they only represent one-fourth of the systems used to narrow down from hundreds to dozens of ideas.

For five-minute CEAs:

- Numbers are mostly either estimated or based on the first result in Google
- Only a few key numbers are calculated within a single row
- Only the largest factors are considered
- Citations are not tracked

Expected outcomes include:

- A quick Google spreadsheet-based CEA that is understandable to the creator
- An endline number that is cross-comparable

Advice for conducting five-minute CEAs:

- Before conducting CEAs we recommend [doing a calibration exercise](#)

- Unless you are using a specific study that reports cost data and has a clear effect size, you will almost always only be able to use intuitions and Fermi calculations
- With regard to effectiveness, think about the theory of change for the intervention, how many people you would need to reach, how many among them would take up the behavior/treatment/ask, and how that will translate into the metric you care about
- With regard to cost, identify the largest factor and think about it in terms of approximate product cost and salary cost for each person reached
- It can be worth spending 1-2 hours preparing a CEA template with key parameters.
 - For example, for family planning research, useful parameters include the efficacy rates of various contraceptives based on a quick literature review. If a large number of interventions on the list target specific populations, consider looking at their share within the larger population and how many people you can reach at different levels (a clinic, a school, etc).
- Consider using an alarm to respect the five-minute limit, as it is easy to go over time with this methodology
 - Try not to use more time, in order to ensure consistency in your mental model for all the interventions. Since a lot of factors might require on-the-spot judgment calls, it is useful to make sure these remain constant across interventions.
- When all the CEAs have been completed, it is a good idea to have another researcher look over them to see if anything seems counterintuitive
 - Results that have extreme values are often due to error, so it's helpful to have a second pair of eyes.
- Different methods result in different outcomes. Considering multiple methods reduces the likelihood that a promising intervention is missed.

Two-hour CEAs

The goal of two-hour CEAs is to compare multiple charity ideas within a given area and to sort them in approximate order of promise before conducting a deeper report. We used two-hour sorting CEAs in one of the four cause areas we considered (mental health) to later allow comparisons to other possible two-hour methodologies that could be used to sort.

Biggest differences between a two-hour CEA and a five-minute CEA:

- Key numbers are based on sources instead of intuition
- More factors are considered
- External CEAs are searched for and values from these are used as inputs
- Citations are tracked (but are not organized for readability)

Expected outcomes include:

- A Google spreadsheet-based CEA that is understandable by the research team
- An endline number that can be used to compare the idea to others and is somewhat reliable

Advice for conducting two-hour CEAs:

- Even within two hours you may fail to find certain relevant data, you may discover that there isn't sufficient evidence, or you may have too low confidence in the evidence you find to create a robust CEA. If this is the case, it is best practice to use either your prior estimate for cost-effectiveness or the average of the top 30 interventions. This can be done using a Bayesian approach and adjusting towards these priors as explained by [GiveWell](#).

20-hour CEAs

The goal of the 20-hour CEA is to provide one of the four major inputs to our conclusion regarding whether it is worth founding a charity based on the idea being evaluated. Twenty-hour CEAs are done fairly late in our research process and are not begun until 36 hours have already been put into other approaches on a given charity idea. At this point, there is a fairly strong understanding of the idea and in the case of mental health interventions a small CEA to build on, created at the previous stage of research. The 20-hour CEA will use all the elements in this document.

Elements that are unique characteristics of the 20-hour CEA include:

- More factors are considered and at a greater depth
 - Counterfactuals are considered formally
 - A higher percentage of the numbers in the model are based on multiple citations
 - Deeper consideration is used when considering the strength of each source
 - External CEAs are used as a minor input

- There is greater readability
 - Endlines are given in both intuitive and unintuitive metrics
 - Intuitive and readable color coding is used
 - Citations are polished
 - The model includes the relevant citations for each number in Google Sheets and Guesstimate without explaining why these sources were used or what the research behind these numbers looks like
 - Assumptions and possible issues are stated in another document
 - This document is where the numbers used are explained, including why certain sources were considered and/or what the research behind the numbers looks like
- More things are double-checked
 - Sensitivity analyses are double-checked
 - A second model is created in Guesstimate to error-check
 - External checking is done

How a 20-hour CEA compares to other methods used in an 80-hour report

- [10 hours] Broad undirected reading and **crucial considerations** (informed consideration)
- [16 hours] Directed research (weighted factor model)
- [10 hours] Finding and talking to experts (experts)
- [20 hours] CEA creation (CEA)
- [4 hours] Directed research (weighted factor model)
- [10 hours] **Summary writing** and internal contemplation (informed consideration)
- [10 hours] Showing endline report to experts (experts)

Expected outcomes of the 20-hour CEA include:

- One cross-comparable spreadsheet + Guesstimate model duplication, both of which are readable and publishable to the general public
 - A Google Sheets model, made before the Guesstimate model (Google Sheets is more reliable than Guesstimate, so you are less likely to lose all of your numbers in a saving error)
- An endline number we can use to recommend the intervention

External CEAs

External CEAs have much diversity in both quality and formatting. Thus, they have a wide range of possible uses. Due to this diversity, external CEAs are basically never directly comparable to CEAs created by other organizations. We see CEAs at roughly three different levels: informative, suggestive, and predictive.

Informative CEAs: Many CEAs, even of low quality, can be informative to generate ideas or obtain citations for key numbers. For this level of CEA we do not take the endline as informative or even suggestive of the intervention's impact, but if we are already investigating an area, we will consider the variables and citations used in informative CEAs when creating our own CEAs. Often, quick or back-of-the-envelope calculations fall into this category.

Suggestive CEAs: Many CEAs have some quality but are not assessing the same metrics we are, or they are not built to apply to the same situation. We see these CEAs often as suggestive that a charity idea could be promising. Suggestive CEAs provide us sufficient evidence that our views often update based on the results of the CEA. For example, we view the [DCP3 CEAs](#) as suggestive, meaning that if an intervention looks cost-effective on their models, that makes us think that related charities and intervention areas could be promising. We do not take the endline numbers literally or even as comparable. For example, if DCP3 says intervention A is better than B but they are both cost-effective, we would do our own comparative research.

Predictive CEAs: Some CEAs are sufficiently high quality or close to the methodology and endlines we consider that they can be taken as predictive. If a predictive CEA were done on an intervention we are considering, we would often give it considerable weight in our process and potentially use many of the same numbers and inputs when comparing. We would view CEAs like this as useful in predicting which areas are better if the same organization has completed multiple CEAs. CEAs we find predictive include those done by [GiveWell](#).

Reading list

[EA concepts: Cost-effectiveness analysis](#)

[GiveWell: Sequence vs. Cluster \(classic\)](#)

[GiveWell: Our Criteria for Top Charities \(their heuristics\)](#)

[GiveWell: Cost-Effectiveness Overview](#)

GiveWell: GiveWell's Cost-Effectiveness Analyses (past examples)

ACE: How ACE uses CEA

Peter Hurford: How Do EA Orgs Account for Uncertainty in Their Analysis?

Peter Hurford: Five Ways to Handle Flow-Through Effects

GiveWell: Guide to GiveWell CEAs

Eva Vivalt: How Much Can We Generalize from Impact Evaluations?

Christian Smith: The Optimizer's Curse & Wrong-Way Reductions

Conclusion & Summary: Part 2

You now have been given a set of tools that can help you make better decisions, but the journey is far from over. The best decision-makers sharpen their tools, practice with them often, and replace the ones that don't work well anymore. Many of the fields described in this writing are rapidly growing or changing and will have new tools and heuristics over time. We should always strive to be better. But what does making consistently good decisions look like? How would we even know what is better? People are bad at judging how good they are at decision-making. As with driving, most people think they are better than average. What are the key principles to remember out of this book and content?

What does good decision-making look like?

Making a good decision often feels the same as making a bad decision to the person making it. However, that does not mean there are no ways to tell the difference.

There are three key ways to get a sense of your decision-making ability:

- Making predictions
- External review
- Results

Making predictions

Making predictions is one of the fastest ways to tell, as there are lots of areas where predictions are quickly tested and proven to be false or not. Think you have found a way to beat the stock market? Try investing a small amount of money in that system and review the outcomes of your predictions. Predictions can also be made on simple things such as the odds of something happening discussed at informal conversations with friends. Try putting real estimated numbers on these guesses and see how often you are right. If your friends stop wanting to play due to you winning too often, that is a good sign.

You can even make predictions absent any other actors. For example, if you think there are two possible outcomes to a given event, try writing down an explicit probability and then review it. You will likely find patterns that differ between reality and your predictions. Most people are overconfident and most people tend to overpredict in a consistently optimistic or pessimistic way. Making predictions gives

you a consistent source of real-world data; just make sure to track them consistently so you notice both the hits and misses.

External review

You might not feel any different when making a good call vs a bad one, but external actors sure will. If you have external people you trust as good thinkers and good decision-makers, asking them to review both your process and conclusions can be a good way to get a sense of your own abilities. Many of the more careful thinkers will be able to recognize a solid reasoning process even if they disagree with the conclusion.

Having an external group of decision-makers you trust review your work also gives ample opportunity for feedback and further growth of your process. Keep in mind this does not mean all these external reviewers will agree with all of your processes and conclusions, even the ones that are right. Even your external decision-makers can be biased, wrong or under-informed. However, they will – on average – be a better judge of where you have gone astray. And heeding their advice strongly for issues you cannot rapidly test is a wise idea.

Reviewing other people's work and getting a sense of their process will also be well worth the time for the benefits it provides to your own process.

Results

Ultimately, good decision-makers will recognize this fact in their results. Good decision-makers won't win every decision, and will be affected by luck (good and bad). However, like a better poker player, they will do better with the cards they are given than others will.

At the end of many decisions, good decision-makers will often be more successful at accomplishing their goals (whatever those goals might be). This is a longer feedback loop as it often takes considerable time to progress on the most important goals, yet ultimately this is the key metric for good decision-making. Can you make decisions that maximize the outcomes you want? In the case of good decision-makers, their charities will be more successful, their relationships will be more stable, and their habits will be more established.

Weakness in one of these areas does not necessarily turn someone into a bad decision-maker. Yet ultimately one of the easiest ways to tell someone's decision-making abilities is looking at their luck-controlled results. Controlling for

luck is important. For many of the most successful people in life, luck played a large role. Keep in mind a lottery ticket winner is not a better decision-maker than the average person (in fact they are likely worse than non-lottery ticket buyers) but the person who wisely picks a career that consistently maximizes their goals of impact and happiness very well would be.

What are the key principles to remember out of this book and content?

This is a long book filled with a considerable amount of information. Research consistently shows that people only recall and use a fraction of the content they encounter. So what are their key principles to remember? Below is a one-page cheat sheet to help remember the most important lessons of decision-making applied to charity entrepreneurship.

Frame your decision well before working on it and sharpen your tools – Toolbox

- Determine upfront how much time to put into activities and time cap based on the importance.
- Set up reevaluation points at a specific date to reconsider your activities.
- If you are narrowing down from many options, set up rounds of iterative depth, with time caps for each.
- Predetermine how many tools to use. Err on the side of more tools instead of fewer.
- Spend the most time on the key choices: co-founder, country, and intervention selection.
- Set up feedback loops as early as possible and try to get feedback as quickly as possible
- Get feedback from trusted advisors; always listen even if you decide not to change your actions.
- Keep learning and practicing with your tools. Be aware you will get better over time.
- Be ready for many of your decisions to be made with far less than 100% confidence.

Understand and utilize many different tools:

- Multi-factor decision-making – The Swiss Army knife
 - Create spreadsheets with different weighted factors to compare complex options.
- Rationality – The hammer

- Think skeptically and narrow down to the crucial considerations of the decision.
- Science – The screwdriver
 - Use evidence as high up on the pyramid as possible and set up falsifiability in your decisions.
- Effective altruism – The wrench
 - Consider the counterfactuals and the limiting factors that could affect your impact.
- Long term planning – The pencil
 - Create a focused theory of change and have a possible endgame.
- Problem solving – The duct tape
 - Generate multiple solutions, compare them evaluatively, and try to solve problems upstream in the future.
- Measurement and evaluation – The tape measure
 - Pick your measured metrics carefully and set up systems to measure them often.
- Task planning – The clamps
 - Have a task management system and focus on the highest impact tasks.
- Independent experts – The cellphone
 - Be careful whose advice you believe, but consider many who supply it.
- Budgeting and finance planning – The saw
 - Run cost-effectively from day one, expand thoughtfully, and commit carefully.
- Cost-effective analysis – The magnifying glass
 - Create detailed CEAs, populate them with strong sources of data.

12 months in: 10 concrete questions to see if you are on the right track founding a charity

1. Are you focusing on a single very impactful thing or have you spread out between many areas, some of which are less effective?
2. Have you progressed fast enough to have started taking actions in your implementation country and spent some considerable time there (months)?
3. Do you have a group of established trusted external advisors such that you are getting intelligent feedback at least once a month?

4. Have you progressed on your one-year plan and completed at least half of it?
5. Have you spent less than 30% of your time on things that are not your primary objective?
6. Do you have an established system to track an important endline metric?
7. Are the co-founders of the project and the leadership team stable?
8. Are you making better decisions for your project than you were a year ago?
9. Do you have a long document of lessons learned?
10. Have you gotten fast and effective at using three or more tools during decision-making?