

Summary

Charity Science has been running for 2.5 years and conducted six experiments (with three failures, two successes and one pending). It has raised an estimated \$300,000 CAD spending about \$80,000 CAD in operating costs for GiveWell recommended charities. Our best guess for a ratio is 1:4, although we include many different ways of estimating impact in our report. We feel as though our biggest mistakes include running too cheaply, not hiring soon enough and relying too heavily on skilled EAs. Our plans moving forward are for half of our team to continue conducting experiments and scale up successes in [Charity Science](#) and for half of our team to move onto the new project of [Charity Entrepreneurship](#).

What this blog post includes

After 2.5 years of working on Charity Science, we felt like it was time to write an internal review. This post aims to be a broad overview of the work Charity Science has done historically, what we've accomplished, to give insight as to what hasn't worked, and to provide some of the rationale and context for our future plans.

[Core concept of Charity Science](#)

[History and summary of experiments](#)

[1\) Grants](#)

[2\) Networking](#)

[3\) P2P events](#)

[4\) Skeptics niche marketing](#)

[5\) Legacy fundraising](#)

[6\) Online ads](#)

[Results](#)

[Expected value](#)

[Operational benefits](#)

[Soft Benefits](#)

[Key mistakes](#)

[Idea Mistakes](#)

[Idea issues we are very uncertain on](#)

[Execution mistakes](#)

[Execution issues we are very uncertain on](#)

[Our plan moving forward](#)

[New high level project](#)

[New experiments](#)

[Long term goals](#)

Core concept of Charity Science

The reason we started Charity Science was because we thought that there was a strong case that we could increase our impact if we fundraised more money for the most effective charities than we spent. There was considerable interest in an organization trying this in the effective altruism community and it also seemed like it was possible it would generate considerable learning value for the community. You can see more information on our core concept and key values [here](#).

History and summary of experiments

Historically we have run six fundraising experiments, three of these failed to raise enough money to justify their time and were thus considered failures. Two of these were successful enough to continue putting time in and scaling up, and one experiment is still pending. Short summaries with key information on each of these experiments is below. Our main success metric is counterfactual money moved, which is basically money that would not have been donated to effective charities without Charity Sciences influence.

1) Grants

Description: writing grant applications to large private foundations and governmental bodies requesting funding for GiveWell-recommended charities.

Counterfactual money moved: we received one grant for \$10,000 USD and possibly influenced one grant for \$10,000+. We also acquired Google Adwords grants acquired for several organizations. We didn't count Google Adwords grants or the influenced grant as money moved.

Time spent: ~12 full-time staff months.

Further reading: our [more detailed review](#) which links to all we have written about this experiment. We considered this experiment a failure.

2) Networking

Description: attending multiple conferences, events and social gatherings of people potentially interested in effective charities.

Counterfactual money moved: ~\$1,000 - ~\$10,000 depending largely on whether effects of networking on peer-to-peer (P2P) fundraising are counted here or in the P2P experiment.

Time spent: ~7 full-time staff months

Further reading: our [more detailed review](#). We considered this experiment a failure.

3) P2P events

Description: encouraging and supporting interested individuals to run a personal fundraising event raising money from their friends, family and colleagues.

Counterfactual money moved: About \$200,000 spread over 4 campaigns of Christmas 2014 (~\$88,000), birthdays (~\$25,000), (~\$25,000) Living on Less and the (~\$12,000) Charity Science walk. Christmas 2015 (~\$50,000)

Time spent: ~12 full-time staff months

Further reading: our [more detailed review](#). We considered this experiment a success.

4) Skeptics niche marketing

Description: Targeting a specific niche with a custom made website and targeted outreach. Skeptics were chosen due to their interest in effective charities and the number of skeptics who have become Effective Altruists.

Counterfactual money moved: ~\$0

Time spent: ~6 full-time staff months

Future reading: our plan [here](#), internal write up [here](#). We considered this experiment a failure.

5) Legacy fundraising

Description: promoting and providing estate planning services for those interested in bequeathing to effective charities.

Counterfactual money moved: Our money moved is an expected value as it's based on planned donations as opposed to executed estates. Thus, even though the original sum is hundreds of thousands of dollars, it has been [discounted based](#) on time discounts, pessimistic attrition rates, probate rates and counterfactuals. Non-time discounted (but adjusted for attrition rates, probate rates and counterfactuals) our estimates are 7k-101k. Add an aggressive discount rate of 50% reduction in value per every 10 years past before the money is donated. This changes our estimates for money moved to 3k - 35k. Overall, we are very pleased with these numbers as our service was only up for an extremely short time and was not aggressively marketed yet.

Time spent: ~5 full time staff months

Further reading: We have not written extensively about this experiment yet and this experiment is still running. We considered this experiment a success.

6) Online ads

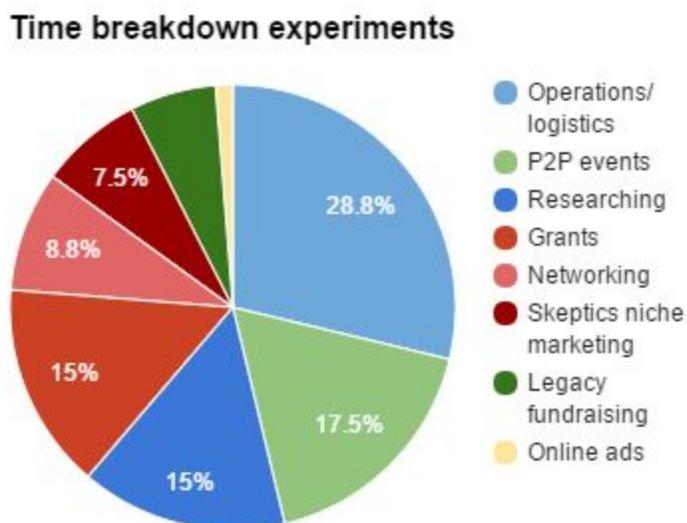
Description: Using free Google Adwords to promote our fundraising and outreach activities.

Counterfactual money moved: ~\$500

Time spent: ~1 full time staff month (although we plan to spend more in the future)

Future reading: Our plan is [here](#). We have not written extensively about this experiment because it is still in progress.

We consider our P2P and legacy experiments successful (green) and our grants, networking and skeptics niche marketing experiments unsuccessful (red). Online ads is still currently being experimented with (yellow). We also spent a large amount of time on operations and research activities (blue). It is worth noting we are not sure which campaign to attribute a considerable amount of money that has been moved through our website. Some of these campaigns may be more successful than our estimates would suggest.



Results

Expected value

We generally view historical ratios as the most reliable metric of success. These numbers are

less subjective than the other numbers that can be used to estimate success. We estimate that after about 2.5 years of Charity Science running we will have raised ~\$300k.

It was recently suggested that we also should estimate the long term fundraising effects of our activities. One way to look at impact is to ask how much impact Charity Science would have had up to this point if we immediately withdrew all founder support and staff. For example, P2P events or our legacy service will likely continue for many years with some positive money moved even if we were to close Charity Science tomorrow. We would guess that our yearly returns if Charity Science was run by volunteers would be close to an average of \$100k raised per year for 3 years and 50k for another 3 years than 10k for an additional 3 years beyond that. In these estimates we have attempted to account for value drift, donor fatigue and decreased counterfactual impact. This totals to around \$480k of expected value over the next ten years. To come to this number, we used much more pessimistic estimates than we have seen other organizations use, but feel as though these numbers are realistic. We are unsure of how to weight this number although are confident we would weight it lower than our historical money-moved numbers. If this number was taken at face value, it could be our expected value is \$780,000 (taking the lower bound from our historical numbers and our pessimistic assumptions on future effects).

We have spent just over \$80k in the last 2.5 years of direct staff and program costs. Our full budget details can be seen [here](#).

A second factor, our community has expressed an interest in staff counterfactuals. Although the specific details of staff's backup plans are confidential, the co-founder estimates of the total counterfactual impact of Charity Science employees/volunteers are close to ~\$500k at the upper bound for all current combined staff/volunteer counterfactuals over the last 2.5 years. This number is extremely soft and based on limited evidence. Using this estimate one could up the counterfactuals included costs of Charity Science to \$580k

We have only listed the pessimistic assumptions, bounds and estimates throughout this report as we find the most optimistic numbers listed tend to be the ones people refer to in conversations and when considering impact, regardless of the qualifications added beforehand. We also consistently find our way of measuring and looking at impact to be consistently more skeptical and pessimistic than other ways of counting and want to be cautious of the urge to post the most flattering numbers possible.

Specific ratios vary hugely depending on what factors to include, how to weight each factor, and how skeptical one tends to be of long term impact. Using the numbers listed above the range can be between 1:11 Charity Science lifetime returns to 1:0.5 Charity Science lifetime returns. For fun we also tried to use identical discount numbers similar to some other calculation methods we have seen and we got a ratio of 1:66-1:128. Our ratios can also change considerably if isolated by experiment (ranging from 1:16 to 1:0). Taking all factors into account yields a ratio of 1:1.3 counterfactually adjusted dollars per dollar spent. Taking only harder,

more evidence-based factors into account yields a ratio of about 1:4. We feel as though these estimates are realistic expectations of what one would expect in the fundraising community. We would be concerned that our measurements might be off if our ratios came back as several times higher than what other organizations report ([page 7](#)). This follows from the common sense rule that extraordinary claims require extraordinary evidence. Finally, these ratios do not take into account operational benefits or soft benefits, which are described more in detail below.

Operational benefits

We have found there have been some organizational benefits that we feel are significant enough to be worth mentioning in this report. The first, and most clearly beneficial one, is the tax benefits we have been able to offer. Canadian donors can now support all of Givewell's recommended charities while receiving tax benefits without being charged the Tides fees or other similar services charge. This has been possible by having a registered Canadian foundation and the extremely helpful collaboration of EA groups like [EACH](#), [GWWC](#) and the Commonwealth market. We do not have an estimate of how much additional money was donated to effective charities because of this benefit, or if there were other benefits to donors, like getting to select from more charities. We would expect this is seen as a large benefit by some.

Registering as a foundation also has other benefits like allowing other projects to grow out from under the same tax status and organizational brand. We feel as though this will aid our [upcoming project](#) and possibly other projects in the future.

Soft Benefits

The final benefits are categorized as "soft benefits". Although we are unsure how to weigh these, many of our donors and advisors think these are some of the largest benefits of Charity Science.

The first of these is personal learning value. Skills like fundraising, marketing, communication and outreach are all very useful and somewhat less common in the more evidence-based communities of charity. We feel our personal skills have massively improved, and our confidence and ability to raise funds for charities is vastly stronger than when we started. We do feel as though these personal benefits have started to taper off as we have spent more time in the field. Another set of skills we have developed are general entrepreneurship and organizational skills. We think these skills have improved and will continue to improve for several years. They should prove valuable in almost any field, but particularly in the fields that the co-founders [plan on going into](#).

The second benefit is communal learning value. We have made great efforts to be [extremely transparent in our organizational activities](#). A big part of why is because we want to allow others to learn from our experiences, mistakes and strategies so that the larger effective charity community can benefit. We have seen various softwares, habits and approaches spread throughout this community, although it is very difficult to draw clear causal lines back to Charity Science.

The final and perhaps largest benefit (in my view) is the team we have built. Over time Charity Science has grown as an organization and now has enough capacity to have staff running two large scale projects (more on our new project at the end of this report). We think the benefits of having this large and skilled team will have significant and far-reaching positive future effects.

It's difficult to weigh all these benefits among more clearly quantifiable ones, but overall the co-founders feel that founding Charity Science was one of the highest impact things they could do with their time.

Key mistakes

Every organization makes mistakes and in our experience the best organizations admit these mistakes publicly. This is not intended to be an exhaustive review of every mistake we've made, but rather what our internal staff sees as the most significant ones we have made historically and the largest considerations we are unsure we are making the right call on. Our self-criticism broadly falls into the two categories of "conceptual" and "execution". These criticisms were initially made to assist an external reviewer who was going to review Charity Science but then became too busy.

Conceptual Mistakes

The first conceptual mistake was the assumption that fundraising on behalf of direct-level charities is preferable to charities simply fundraising for themselves. We have found that fundraising is more difficult if you are not the direct charity. There are a number of possible reasons for this, including:

1. The fundraising methods available to external groups like us are much more limited than for the effective charity itself. For example many grants we applied for needed specific and detailed information about the founder that was hard to give as an external organization.
2. People get attached to the people they are interacting with personally so if they interact with Charity Science staff, they tended to like and want to fund Charity Science instead of the charities we recommend.

3. There are added logistical complexities of being a separate organization such as having to do separate budgeting, donation receipting and re-granting. Additionally our model was confusing to many people which made fundraising from them harder.

The second mistake was the belief that science could be easily and effectively applied to fundraising. We consistently found the [research weaker](#) and less clear cut than we were expecting. Even when talking to experts or paid consultants, the amount of empirical data or clear expert opinion was much weaker than expected. We found we gained much more knowledge from executing small-scale experiments than we did from online research or talking to experts about different strategies.

Conceptual issues we are very uncertain on

One issue is the possibility that better research is more likely to bring in more funding than explicitly focusing on fundraising or outreach. This is somewhat supported by the much larger amount of funding that an organization like GW has moved (compared to Charity Science). We are currently very uncertain about whether this is true or not and can see a strong case for both viewpoints.

Another issue is that fundraising may be better done in a much broader way. For instance, moving people from local charities to global charities instead of from one specific charity to another. This might be more effective or more persuasive to people. We feel as though this comes down to speculative guesses regarding both the marketability of broad versus specific approaches and comparisons between effective charities and global or local ones. We do not have confident estimates on either of these issues.

Fundraising may be far less effective than standard money-moved ratios might make it seem because the difference between the charities that people switch between may not be that large (e.g. Oxfam to AMF). This is another issue we are uncertain about, although more certain than the above two. This might mean earning to give might be better even if it's a smaller amount of money moved per person, it is money that otherwise would not go to any charity.

Finally, possibly our largest concern is how our time would be best spent counterfactually. We are extremely uncertain this is the case (although we were more certain when we founded Charity Science). It seems plausible that it would be better for highly dedicated EA's to work on founding and improving charities with the intention of having governments or very large funders cover the costs. This is particularly the case with large effective foundations such as Good Ventures and the Bill and Melinda Gates Foundation, who seem like they would be keen to fund extremely effective charities. That being said we do think moving more money to great charities (via earning to give or fundraising) is a great fit for most EAs, and it's rare to have the right dedication level and skill set for the former option to be a good fit.

Execution mistakes

Our first mistake was running our organisation too economically. We spent too much time trying to save money in ways that, in hindsight, were not an effective use of our time. We think this was somewhat a reaction to other organizations we had seen making the opposite mistake. This might also have been tied to our focus on ratios instead of net money moved.

Our second large mistake was not hiring soon enough. Because we wanted to run at a low budget we were averse to hiring staff. However, we have found hiring staff to be hugely beneficial and overall has made our organization much more effective. We likely could have run all of our current experiments over half the time if we had hired aggressively from day one.

Our third mistake is relying on the generally high skill of the co-founders and other highly involved EAs. We think that Charity Science likely could have been run by a single highly involved EA and a team of people new to EA concepts or employees from more standard nonprofit circles. This choice was likely affected by our worries about cost/budget and our underweighting of staff counterfactuals.

Our fourth mistake was an underutilization of more traditional business practices (e.g. clear hierarchies, more specialization, clearer delegation systems). We feel this is somewhat inevitable given the youth of our charity, but we still feel we could have executed better by taking a more traditional approach before experimenting in these areas.

Execution issues we are very uncertain on

We have some concerns about counting the counterfactuals of donations from EAs. We have tried extremely hard to be self-critical and in some people's view go too far in undercounting our impact. We have thought a lot about this process, but are still not fully confident whether we overcount or undercount our own impact. ([An example of how we do this](#))

Prior to Charity Science, our team had a low level of previous fundraising and outreach experience. We are unsure about the difference this would have made, as our advisors and consultants did not have nearly as much insight as we were expecting. Even the consultants with impressive track records (e.g. they say they have raised millions a year) seemed much less impressive when we dug deeper—often they only raised a few hundred thousand marginal counterfactual funds and the rest is just a build up of the donor base over the decades of existence of the charity.

We are unsure if we focused too much on quick output or measurable activities. Currently our experiments have succeeded at quite a high rate (33%) but we have not had any exceptionally

large wins (e.g. a 5 million dollar grant). It's possible that taking a riskier approach to fundraising would yield more returns in the long run.

Our plan moving forward

New high level project

The biggest change to Charity Science is that 50% of our current team (including the majority of co-founder time) is moving to the new project of [Charity Entrepreneurship](#). We feel that this is an extremely high impact project and allows us to move closer to the ideal model of Charity Science (with less reliance on founders and highly experienced and involved EAs). This will largely lower the counterfactual cost of Charity Science in the future. This experimental project may also answer some of the uncertainties we have around fundraising as an external organization versus being the direct charity. Our full plans for Charity Entrepreneurship can be found [here](#) and some more rationale to why this might be a high impact project can be found on the [EA forum, GiveWell and 80k](#).

New experiments

Over the next 6 months we have a new batch of experiments that we are considering with a stronger focus on breaking out of the EA movement and reaching new communities. We expect to run 2-4 experiments over the year as well as maintaining past successes. These new experiments are face-to-face fundraising, further online ads testing, donor stewardship and helping volunteers run small scale experiments. You can see all of the areas we are considering [here](#).

Long term goals

Our long term goals are to continue to experiment for approximately the next year and to run successful experiments indefinitely. We expect to hit diminishing returns on fundraising experiments somewhere in 2016 or 2017 and will change our major focus to scaling up success or scaling down the required staff time. We expect to lower Charity Science's counterfactual cost substantially in 2016 and to raise its net money moved slowly over time. Our plan is for Charity Science to take minimal co-founder time while continuing to generate passive impact.

We see four possible very long term plans for Charity Science a [growth plan](#), [a stable plan](#), [volunteer based plan](#) and an integration (into a larger direct charity) based plan. These each

have pros and cons and we feel like we will be better able to evaluate these options middle or end of 2016.