

Developing World IV: GiveDirectly


Professor Paul Niehaus
University of California, San Diego
Department of Economics

05/01/2018

These notes are based off a presentation by Professor Paul Niehaus for the section on the developing world in the Mechanism Design for Social Good Reading Group. The notes are taken by members of the reading group with some figures and texts taken from presentation and website. Questions and comments from reading group members during the presentation are labeled as such. Please contact the organizers with any questions or comments.

1 Introduction

Professor Niehaus co-founded GiveDirectly as a graduate student. It was inspired by the idea that the end of extreme poverty is “in reach.” We would need to spend about 65 billion USD a year to “end” extreme poverty. Foreign aid is also much greater.



This raises the question: *why are we not closing the poverty gap?*

In 2000, we started doing development economic experiments. We have *a lot* of rigorous evidence on the effectiveness of cash transfers:

- People generally use money responsibly: we do not see reduced labor supply and if anything see lower spending on alcohol and tobacco,
- Cash transfers have a wide range of positive impacts, including long-term ones in health, labor supply, etc

- Short-term transfer generate long-term impact for at least some. Understanding who these people are and how to identify them is a new research frontier.

There have been significant advances in last-mile payment technologies (e.g, through mobile phones). This makes remote, large-scale, secure cash transfers feasible. Mobile money has been a dominant model in many African nations.

Western aid is at least 90% in kind. Humanitarian is at least 90% and in development aid and private charity cash is not systematically measured, but we believe are negligible. This is even higher for some countries. e.g, 100% of US aid is in-kind and there is no cash aid. The people being helped have almost no say about the form of aid. The entire system is built on the notion that intermediaries should decide *what* the poor receive. Humanitarian aid is designed around categories (food, health, education, etc) of specific goods. Flexible cash transfers have played little to no role.

The amount of money that is actually being received by the poor is not transparent in most aid settings. Often, the rate is very low, and most is taken up by overhead. And, there is no measure of value, let alone mechanical amount of eventual transfer.

2 Overview of GiveDirectly

GiveDirectly is championing a much simpler approach. It manages transfers end-to-end using electronic monitoring and payment technology. (Other charities are more front-end for the donors.)

The steps are:

1. Program design: targeting, payment frequency, payment amount, nudges, etc
2. Enrollment: conducted in-house, provide users SIM cards, enroll participants in money money.
3. Verification: eligibility, independent audits, automatic checks, etc
4. Transfer: fully integrated mobile money system
5. Monitor: experimental impact evaluation. Also call-backs to process feedback (e.g., could you register? Did you get your money? Were you asked for a bribe?) There is a 98 percent follow-up rate
6. Evaluate: RCTs by external academics. There are 8 or 9 running now. Having external researchers is key here. Third parties do data collection (e.g. JPAL, IPA). The plans are public, pre-specified, and audit-able.

We first locate extremely poor communities using publicly available data. We then send field staff door-to-door to digitally collect data on poverty and enroll recipients. We target households using criteria that vary by region including aggregating a range of factors or looking at housing materials.

We use a set of independent checks to verify that recipients are eligible and did not pay bribes, such as physical back-checks, image verification, and data consistency checks. For example, we use GPS coordinates and crowd-sourced labor to detect irregularities. We have experimented with using satellite imagery.

We transfer recipient households roughly \$1,000, or around one year's budget for a typical household. We use electronic payment systems; typically, recipients receive an SMS alert and then collect cash from a mobile money agent in their village or nearest town.

We call each recipient to verify receipt of funds, flag issues, and assess our own customer service. We also staff a hot-line for inbound calls and in some cases staff follow up in person.

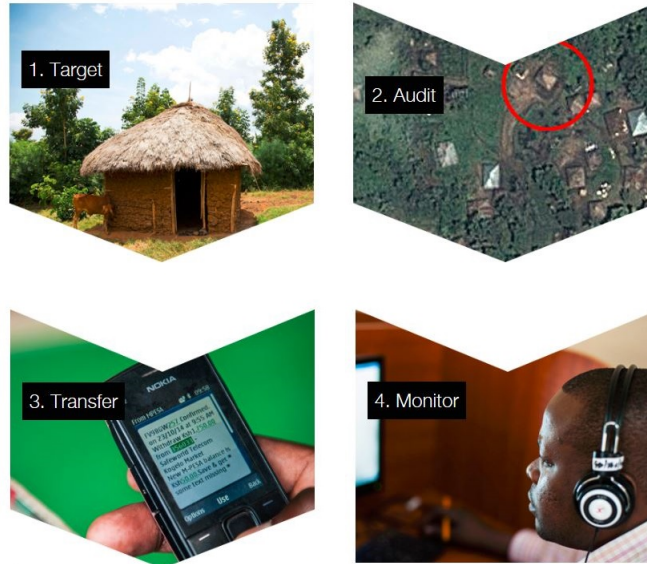


Figure 1: Program model for GiveDirectly, an end-to-end cash transfer program.

Question: You mentioned that intermediaries have downsides. But, all the auditing that you have to do might be more efficiently done by a local organization. What are the downsides of having the end-to-end process done by one organization rather than getting local organizations involved?

Answer: The main downside to intermediaries is having worse bribe control. It is harder to make sure that the right people get the money. There is less information and reduced ability to fire cheaters. Ultimately, you're trusting someone else to disburse the money if you use a local organization. In terms of cost, it's unclear. The main issue is that if you want to grow something quickly, then doing it yourself takes time.

Question: Have the issues you're raising been getting easier over the years?

Answer: The problem does not get easier, but GiveDirectly is getting better at solving the problems. For instance, more complicated targeting rules are more game-able, but we have simplified these rules.

GiveDirectly is expanding into other countries. We have grown local staff. There is a small team in New York doing core infrastructure and marketing in the U.S. Historically, 89 to 90 percent of each donated dollar ends up in the hands of each recipient. The team is on a 50 million USD per year scale. This is comparable to major public sector programs in these countries. Most countries have unconditional cash transfer programs. donors are the laggards. Host countries have been early adopters of cash transfers.

People who like GiveDirectly tend to be in the Effective Altruism, tech, and finance community. The longer-run goal is to reach out to a broader group of people who are potentially less analytical about how they give. The press coverage has changed from "these guys are crazy" to "why aren't we doing this more?" We have consistently been rated among the most impactful ways to give and have driven public debate on the future of foreign aid. But, money hasn't followed yet.

To be clear, we don't think all aid should be cash. But, a lot more of it should be. We should think of things like an index fund for development spending. Cash transfers are like a low-cost, passively-managed approach. Aid administrators are like fund managers. They should have to argue that their costs are justified by the excess returns you can generate. Right now, no one is

under any pressure to make that argument, but, in many cases, they could. There are many issues with large externalities, public goods problems, vaccine research and development, where there's clearly a role for government or external parties to regular/subsidies. What we'd like to see is pressure for people to make that case in an evidence-based way.

The sociology of index funds is also very interesting. People were mocked at first, but in the end they won. Cash transfers are a fundamental challenge to the aid industry model, and there is similar social resistance.

Last thing they're working on is the human side. They have a live-stream of recipient feedback from GiveDirectly. It is unfiltered, and not just the curated stuff most charities send out. The goal is to humanize and be transparent.

Giving someone money is an adventure; it's like investing in an entrepreneur. If you look at the feedback and data, you can see all sorts of expenses. Some people purchased live-stock, invested in land or farming equipment, started small business, bought cars/motorcycles, purchased music instruments, etc.

3 Discussion

Question: How do you decide who gets what?

Answer: The broad area related to this is "poverty targeting." How do you determine who is poor? There are two main approaches in developing countries:

- Community-based targeting, where the community ranks members in terms of poverty, then these reports are used to allocate resources, and
- Proxy-means testing, where information is collected based on poverty indicators. This is used to determine how the community ranks members.

The second is prominent, especially for public sector programs. We have experimented with many forms of both, and we have found that this issue should be simpler. We started sending money to people with homes with thatched roofs and mud walls. (This could be verified using satellite imagery.) That was already simple. Now, we just use administrative data to find poor villages and enroll everyone in that village. The added cost and complexity to find a finer-grained metric is not worth it. Furthermore, even the metrics that were previously used are not necessarily a great predictor of economic welfare. A complicated formula is a nightmare to implement. It presents more opportunities for discretion, gaming, and fraud. Universal Basic Income are at the extreme end of the spectrum in terms of complexity, and we are experimentally testing that right now in Kenya.

Question: You mentioned that there is new work on how to predict long-term benefits from one-time cash transfers. What have people found? How would you work with what you are suggesting?

Answer: You have identified an important tension. My personal preference would be to redefine the objective away from narrow goals like "health" to the broad goal of "whatever recipient prioritizes" but also realize that pragmatically it may make sense to figure out how to do cash transfers in a way that maximizes health impacts so that health funders can get on-board.

Question: What issues come up when you are trying to scale this up? You mentioned that one benefit of GiveDirectly is that everything is done in-house. Does that affect scalability?

Answer: The main constraint has been on the revenue side. Basic operational models haven't fundamentally changes. There are already government programs doing cash transfers on a much larger scale. The issue is not primarily scalability on the operations, although that is a challenge.

The fact that GiveDirectly does everything in house is a constraint on scalability. The main issue is getting donors to fund it at a large scale. Scalability of persuasion is an issue. In terms of thinking about building versus outsourcing, building things ourselves tends to be more scalable than outsourcing. (This is not a hard principle, but it has been the case for us.)

Question: Is it possible that scalability and persuasion are related to one another? Can you inform the donor when you have more granular data?

Answer: Yes, you need large scale experimentation to get such detailed effects.

Question: With respect to scalability of persuasion, have you been able to experiment with what gets people to donate?

Answer: There is a lot of research on this, and more work to be done. GiveDirectly is at the beginning stages. I have mostly been involved in experiments with the larger institutional funders, and what it takes to get them on board. We had a partnership with USAID. They have a procurement mechanism called the global development alliance. You have to fund at least half of the project from the private sector, which allows you to propose ideas to USAID. Then, we worked with the general counsel's office to come up with rationale to justify using taxpayer dollars on this since USAID is sensitive to not using tax-payer money to things voters do not like. Then, there is the issue of finding the internal people who want to champion this. We began with an economist in Rwanda who is now back in DC coordinating some experiments. But, USAID is doing impact evaluations in a few years, sometimes having head-to-head comparisons of cash transfer programs with their counterpart non-cash transfer programs.

Question: You do a lot of analysis at the village level. It must be hard to evaluate whether a village is more or less poor. Could you reverse engineer things to see which villages benefited most from transfers and what village characteristics seem to predict this?

Answer: This is a wide-open area. All targeting historically was based on finding the deprived people. No one has asked how you target the people for whom the intervention makes the biggest difference. That's a large gaping hole in development. We know more about cash transfer than other things because there's more research, but this is still very wide open. This could be village-level characteristics, individual things (e.g. age and other demographic structures), or something else that we have not yet identified. This is an interesting avenue for research.

Question: Consider an academic study where in one village you give people money directly and in another village you have a program manager decide what to do with the money. This is a difficult comparison since it is sensitive to who the program managers are, finding identical villages, and so on. How would you conduct such a study?

Answer: The key question is the one you asked: if we give the manager the money, and a salary, what would we get? Increasingly, studies have a one treatment arm be in-kind and cash. But that design misses an organizational design question: who decides what to provide? It is something I am trying to do: shift the conversation towards the one you posed. In addition, which outcomes you measure is a tough challenge. If you get a narrow intervention, the outcome may be obvious, but not for cash transfers. Different people use it on different things and this can impact an enormous array of outcomes. What do you need to measure? How do you aggregate it? We've tried to look at a broad set of outcomes and leave it to others to weigh these outcomes. We will be more proactive about providing a framework in the future.

Question: Have there been issues that come up with respect to the gender, age, and other demographic characteristic of the recipients. E.g., in many developing nations, women may not have access to their own phone. Is it challenging to make sure that the cash transfers are getting to all

individuals for whom it is intended?

Answer: This has not had an impact on who we decide to help. We give everyone a SIM card. We have given people the option of purchasing a mobile phone with their cash transfer. Most people take them up on it. Possession is not a constraint. There is a design choice about whether to enroll the husband, wife/wives, etc. We have done experiments comparing this head-to-head and have not found large differences. This could be because of power structure or aligned preferences. In the UBI experiment discussed above, they are enrolling every adult. The one constraint is enrolling young people since you must be at least 18 years of age. There is a strong argument to enrolling people sooner, especially for teenage girls since they will be making important life choices and financial independence could have large, positive impact for this demographic. We have not found out how to get around that constraint yet.

Question: Could you give that amount to their parents?

Answer: We could do that, but the downside is that they would not control their money. We could also try to build on some mechanism to verify that it is being used well, but that gets complicated very quickly.

Another issue is: currently, transfers are individual. There is a major outstanding question for how to finance local public goods. We could be impacting building additional mechanisms for that purpose: e.g., letting people vote on asset creation. This would raise interesting mechanism design questions.

Question: Do you have data on demand for public goods?

Answer: No, but that would be a good place to start. There is a strong sense that local public goods are under-provided. The information we need would be people's willingness to opt in for public goods provision. We did a round of experimentation in Kenya. Individual transfers did not generally lead to increased government revenue or public goods provision. It seems like you would need an additional mechanism or layer of infrastructure, assuming that there is demand.

Question: How can our group get involved?

Answer: There are a range of things that people in this group could work on. Public goods and heterogeneous treatment effects are things I've been thinking about. Ideas are welcome here. Another thing is: there are a couple of principles that come out of the cash transfer work. (1) cash is a benchmark, and (2) cash can be implemented in a specific way. Can you give people a choice? e.g., do you want \$100 or do you want good X? That makes you accountable to them. You must convince the recipients that good X is a good idea for them. Accountability is really important, and a principle that you can apply to many things.