

A conversation with Paul Niehaus and Carolina Toth, September 7, 2015

Participants

- Paul Niehaus – Co-Founder and President, GiveDirectly
- Carolina Toth – Interim Vice President of Fundraising and Operations, GiveDirectly
- Rebecca Raible – Research Analyst, GiveWell
- Natalie Crispin – Senior Research Analyst, GiveWell

Note: These notes were compiled by GiveWell and give an overview of the major points made by Dr. Niehaus and Ms. Toth.

Summary

GiveWell spoke with Dr. Niehaus and Ms. Toth of GiveDirectly for an update on its activities. Conversation topics included potential partnership projects and research studies, effects of these projects on GiveDirectly's core activities, room for more funding, and other programmatic, operational, and financial updates.

Cash transfers measured against other interventions

Cash transfers currently represent an estimated 5% of all humanitarian spending, although some, including the UK's Department for International Development (DFID), believe it should make up a higher percentage. Cash transfers are still a controversial intervention, and disagreement is emerging within the field of humanitarian policy over the degree to which they should be used. People who are invested in maintaining the status quo in international development are likely to push back against cash transfers.

Institutional partnerships and policy work

Impact of partnerships and policy work on GiveDirectly's core activities

GiveDirectly is engaging in or considering a number of partnerships that would involve implementing cash transfers, and also could have policy impact. Attempting to affect policy change is very high-risk, but potentially also very high-reward. It is also much harder to measure the impact of this work than the impact of a direct cash transfer.

GiveDirectly does not believe its policy work is significantly different from its traditional cash transfer campaigns. In its standard activities, GiveDirectly will make decisions about who gets transfers and how recipients are randomized. When GiveDirectly is attempting to affect policy, these decisions are still made, but with explicit policy impact in mind. From this perspective, policy decisions are an extra input that precedes the execution of a project, in order to ensure that the direct action of transferring cash has the greatest possible impact.

Regardless of its policy work, GiveDirectly intends to maintain its current level of cost-effectiveness by delivering 90 or 91 cents of each donated dollar directly to transfer recipients, which remains the minimum impact of the organization's activities. An indirect effect of increased funding for GiveDirectly's cash transfer programming, as well as increased positive media coverage, has been an increase in resources for GiveDirectly to invest in policy work.

Funding structure

For any large institutional partnership, GiveDirectly pursues a 50/50 split of the costs—GiveDirectly or GiveDirectly's funders contributing money to the project increases the probability that a large institutional funder would contribute as well.

Allocation of time by GiveDirectly management

Ms. Toth estimates that Dr. Niehaus spends no more than 50% of his time on institutional partnerships and believes that this is the area of GiveDirectly's work where the marginal value of his time is currently highest. Progress on partnerships with large organizations tends to be slow, as they require approvals from several people.

Criteria for participation in partnerships

To decide whether to participate in a collaborative project, GiveDirectly must weigh benefits against costs. Many partnerships would incur fixed costs of moving to a new location or finding logistical solutions in a new area, as well as other risks and uncertainties. Potential benefits could be changing policy or changing mindset among a group of people. In particular, GiveDirectly considers the counterfactual question of what might happen if it decided not to pursue the project.

For example, in one potential partnership that GiveDirectly ultimately decided not to pursue, the planned study already included a cash transfer arm as a benchmark, and GiveDirectly felt that it would add little value by acting as the implementing partner. GiveDirectly decided to limit its role to providing advice on how to design the transfers.

Another partnership GiveDirectly explored was working with two large multilateral organizations to set up a comparison between cash transfers and cows in one country. GiveDirectly believed that while the potential for impact on government policy and recipient outcomes was high, the risks of moving into an unfamiliar area and using unfamiliar payment systems were too great. It was also unclear whether a large enough sample could be assembled for the cash transfer arm to constitute a scientifically credible study.

Work with DFID

DFID invited members of GiveDirectly's board to participate in a panel discussion about the role of cash transfers in humanitarian aid. However, it is hard to predict what kind of impact on policy might result from the panel's findings.

Benchmarking initiative

GiveDirectly has had conversations with a non-profit that runs evaluations of aid programs about how to promote more head-to-head comparisons between cash transfers and other interventions. An informal poll among lead researchers of the non-profit's ongoing studies indicated a high level of interest in incorporating cash transfer arms. GiveDirectly likely would not undertake the work of setting up cash arms for studies due to the high costs involved, but in many places there are in-country implementing partners that could execute a cash transfer arm if they had the necessary funding.

GiveDirectly would like to learn whether grantmaking organizations would be interested in funding the addition of cash transfer arms to studies. However, it would not itself fund other organizations to do this. GiveDirectly intentionally does not regrant to other organizations because the process of evaluating sub-grantees would require a huge investment of resources that GiveDirectly is not currently willing to dedicate and that would reduce efficiency.

Financial updates

Room for more funding

Field directors' cash transfer capacity

GiveDirectly believes that its field directors in Kenya and Uganda will each have a cash transfer capacity of \$17.25 million in budget year 2016 (this does not include January and February of 2016, which belong to budget year 2015). GiveDirectly is hiring a flexible field director (i.e., one who works on a variety of projects in multiple countries), Eric Friedman, who will start at the end of September.

Future hiring opportunities

Hiring new field directors would not be a challenge for GiveDirectly. If it had more funds to distribute via cash transfers, it would hire more field directors and thus raise its capacity significantly. GiveDirectly's COO, Piali Mukhopadhyay, has the capacity to manage five field directors.

Hiring more junior staff also presents no obstacles, as unemployment in the countries where GiveDirectly works is high. Hiring for these positions might take a month or two of the field directors' time.

Projected use of Good Ventures funding

GiveDirectly now has \$25 million from Good Ventures, as well as some funds raised in the last couple of months, in unbudgeted funds. GiveDirectly would like to fully explore the possibility of finding matching funds for the Good Ventures grant before committing the Good Ventures funds to cash transfers. Once GiveDirectly has a better sense of the likelihood of securing matching funds and the pros and cons of such an arrangement, it will start to consider using the Good Ventures funding unmatched. It is possible, but somewhat unlikely, that GiveDirectly will find a large-

scale matching partnership very quickly. A more realistic outcome is that the search for a matching donor will extend through 2016 and beyond, which means that GiveDirectly would not use any of the Good Ventures funding in 2016.

Effect of inflation on transfer sizes

Background

A few years ago, Kenya experienced a very sharp rise in inflation (25–30% in 2011), which reduced the value of the transfers GiveDirectly distributed.

Until recently, the transfer size had been pegged to an amount in Kenyan shillings (KES) or Ugandan shillings, equal to about \$1,000 based on 2011 exchange rates, and was held constant at that amount. As a result, the transfers in Kenya have decreased in value along with the KES, which has masked GiveDirectly's efficiency gains since 2011, as its costs (such as wages and transportation costs) were increasing due to inflation while the size of the transfers was not.

Adjusting transfer sizes for inflation

GiveDirectly's board has thus decided to raise the amount of the transfers in Kenya by about 30%. Transfer amounts will be indexed for inflation going forward based on regular reevaluations of the inflation rate, which will likely take place every six months when the budget is refreshed. Transfers that have already been committed to recipients will not change; Kenyan recipients who were told in January 2015 that they would receive KES80,000 will still receive that amount. Recipients currently being enrolled, however, will receive about KES110,000. The historical record of GiveDirectly's transfers will not change, as the funds represented by those transfers have already been committed.

In Uganda, the opposite effect has been observed, so that the value of the Uganda transfers is equivalent to about \$850 or \$900.

Reasons for delay on adjusting for inflation

GiveDirectly was aware of the inflation problem in 2011, but did not address the issue at the time. GiveDirectly also mistakenly thought the rate of inflation would be about 2–5%, which was much less than the actual rate. In addition, in 2011, all of GiveDirectly's programs were operating on the stop–start model, so there was more uncertainty around the future of its programs and less focus on long-term trends. The current rolling model in Kenya has brought a shift in mentality, which helped lead to the decision to adjust for inflation.

Policy on foreign bank holdings

To protect against foreign exchange risk, GiveDirectly's policy is to only hold funds within a country if they have been committed to transfers; other budgeted liabilities are currently retained in U.S. bank accounts. However, GiveDirectly is now debating whether to increase the amount it holds in foreign currency to an amount equal to a six-month budget. This would provide a hedge against foreign exchange risk and

some protection against deflation or inflation, but it would also result in greater bank risk. A six-month budget for GiveDirectly is currently \$10–15 million, which is a large sum to entrust to a foreign bank, and GiveDirectly does not yet have an adequate understanding of Kenyan bank risk to make this decision.

Setting efficiency targets for each campaign

In the past, GiveDirectly did not explicitly set efficiency targets for each campaign. Its target was consistently 91%, which represented approximately the efficiency level it had already achieved. GiveDirectly's goal at the time was to maintain that level of efficiency as it built up its program infrastructure. In Uganda, efficiency targets have typically been lower because GiveDirectly's costs in Uganda were higher at the beginning of its operations there.

The most recent efficiency targets were determined during the budgeting process. The person who creates the budget for a country (usually the country's field director) looks at the costs, determines how many field officers are needed and by how much their salaries should be increased, and calculates the target efficiency based on these numbers.

Program updates

Uganda

Funding for October 2015–February 2016 campaign

GiveDirectly plans to begin enrolling recipients for a discrete campaign in Uganda in October and November 2015. The \$4.5 million GiveDirectly currently has is enough to fund enrollment through February 2016; for any enrollments beyond that, more funding would be required. If GiveDirectly has enough funding to transition Uganda to a rolling model, it will incorporate the discrete campaign into a rolling campaign. If it does not get enough funding, it will end its program in Uganda, though GiveDirectly considers this an unlikely outcome.

Enrollment

GiveDirectly is currently trying to secure government approval for all 7.3 million households in Uganda to participate in its cash transfer programs. However, if this does not happen, GiveDirectly would still be able to continue enrolling recipients in Uganda for a couple of years, as Bukedea District alone contains approximately 27,000 households.

Mobile money partnership

GiveDirectly is testing a new partnership with Centenary Bank in Uganda to distribute its transfers, and is also conducting a pilot to test a distributed cashout model with its current partner, mobile provider MTN. A distributed cashout model allows recipients to withdraw their cash from any MTN agent at any time; under a payday model, recipients are told to assemble at a certain place and time to withdraw their cash. GiveDirectly believes that using distributed cashouts where

possible is preferred. In Uganda, MTN is the mobile provider with the broadest distribution network, making it GiveDirectly's preferred partner.

Previously GiveDirectly thought that a distributed cashout model would not be feasible in Uganda, as cell phone penetration is limited and recipients would have to travel long distances to reach mobile agents. This is why programs in Uganda began under a payday model. GiveDirectly is now trying to improve the effectiveness of the distributed cashout model in Uganda by communicating extensively with recipients about where agents are located.

Kenya

Enrollment

GiveDirectly has nearly exhausted its ability to enroll households in Siaya County, although it would be possible to treat the control group in Siaya from the General Equilibrium (GE) study. The control households make up slightly less than half the eligible households in Siaya. However, GiveDirectly wants to maintain the Siaya control group as a control for long-term follow-up purposes, and it does not anticipate problems getting further approvals from the Kenyan government for work elsewhere.

GiveDirectly has county-level approval to enroll approximately 70,000 more households in Kenya, the vast majority in Homa Bay County, but still needs approval from some specific districts before it can begin operating in them. However, previous approval for a district in Homa Bay County took only a few days to secure, so this should not be a barrier going forward.

Rwanda

GiveDirectly will soon begin a \$4 million study comparing cash transfers to other interventions in Rwanda. The study is funded by Google.org and a bilateral aid agency. GiveDirectly's funding partners have provided a list of potential programs against which to benchmark cash transfers. GiveDirectly is now narrowing down the list to the options that make most sense as a basis for comparison. In about two months, GiveDirectly expects to have more details about the study design and which program will be compared to cash transfers.

In the future, GiveDirectly may invest its own funding in Rwanda for other cash transfer campaigns, but this depends largely on the total amount of funding it receives and on how GiveDirectly's work in Rwanda progresses.

Modified criteria for eligibility

The presence of a thatched roof on a household's dwelling will continue to serve as a proxy test for extreme poverty whenever GiveDirectly enters a new area. However, GiveDirectly is exploring other tests to see whether they perform better at targeting households for cash transfers. In Homa Bay County, thatch is not commonly used, so a different proxy means test is needed.

GiveDirectly has tested the performance of potential eligibility metrics by examining the internal consistency of their results, as well as comparing metrics to each other and to the GE study baseline data. Community-based targeting (which relies on reports from other community members about a household's status) has not thus far been a strategy that has performed well. GiveDirectly has found that different groups of respondents give very different answers, or that the same respondents will give very different answers on multiple tests.

Trustee and helper system

To help some recipients obtain their transfers, GiveDirectly has instituted a system of trustees and helpers. A trustee registers for a mobile money or a bank account on behalf of a recipient who cannot do so. This is usually because the recipient does not have a national ID and/or is either very old or too young (as in a child-headed household) to register for a money transfer modality. Trustees are chosen by asking recipient whom they most trust, ideally out of earshot of other family members who might try to take advantage of them. The selection is then validated through discussion with neighbors. GiveDirectly prefers naming trustees who are also transfer recipients, as this reduces any incentive to steal the transfer and GiveDirectly can re-direct the trustee's own transfers if any abuse of the relationship occurs.

A helper is a position of lower responsibility. Helpers may be chosen for recipients who are very old or blind. The transfer can thus be registered in the recipient's name, but the helper retrieves the transfer or assists the recipient in using the phone.

Security risks

Risk of targeting for robbery on cashout days

Large volumes of cash brought into villages on cashout days, making them a possible target for robbery. However, the banks or mobile money providers that GiveDirectly works with provide armored vehicles, security personnel, and valuable cash handling expertise for the cash out days.

Risk due to increased prominence

GiveDirectly does not anticipate increased risk of robbery as it becomes better known in the areas where it works. In Siaya in Kenya, GiveDirectly was extremely well known, so it seems unlikely that criminal activity in new areas where GiveDirectly is less recognized would be greater than it was in Siaya. However, likelihood of crime might vary from region to region. GiveDirectly does not consider security risk to be very high in Homa Bay County, but acknowledges that it could be more of a concern in an urban area.

One of the benefits of the distributed cashout model is that it makes it easier to keep recipients' identities secret and therefore limit the potential for robbery. As part of its follow-up procedure, GiveDirectly asks recipients security-related questions,

such as whether anyone knows that they are recipients and knew when they went to cash out, to monitor security risk.

Policy on reimbursing stolen transfers

When theft of a transfer occurs that is out of GiveDirectly's control, GiveDirectly generally does not reimburse the recipient, although it will try to help recover the money by negotiating with the thief or involving the police or local government.

Impact of research studies on GiveDirectly's activities

Approach to incorporating research studies

A large number of GiveDirectly's enrollees are now participants in one of several ongoing research studies. The model for some of these studies may be significantly different from GiveDirectly's regular campaigns (i.e., the transfer amounts may be higher or lower). When a study's model deviates significantly from the way GiveDirectly publicly presents its campaigns, GiveDirectly finds an independent source of funding and does not use any retail donations to finance the study. For example, one study was funded entirely by the Rockefeller Foundation. The ideas42 study, which also represented a substantial deviation from GiveDirectly's core model, was also independently funded. Examples of retail donation-funded research projects are the GE study and the forthcoming aspirations study, which both use the model described on GiveDirectly's website.

GiveDirectly is now trying to include all of its recipients in either a study or a learning objective. Learning objectives are internal tests conducted by GiveDirectly in which the standard campaign model is slightly varies. GiveDirectly uses results from learning objective campaigns to gather information about how it might adjust its operating procedures (e.g., experimentation with proxy means tests).

Research studies under rolling-campaign model

The rolling model is preferable in most ways to the discrete-campaign model. It allows GiveDirectly to employ people for longer than a few months at a time, and to operate more as an organization rather than a standalone project.

Though it is not necessarily more difficult to incorporate a study into the rolling model, it does introduce some complexities. The rolling model places some pressure on academics to organize their studies in a way that fits into GiveDirectly's timeline. The sooner an academic team can begin its baseline study, the larger the sample it can use. GiveDirectly may occasionally delay some transfers to help preserve a study's integrity, but generally it prefers not to suspend its activities due to delays by researchers.

The GE study in Siaya is an example of how research studies can affect GiveDirectly's timeline. In that case, GiveDirectly moved more quickly than anticipated, and so had to delay the token transfers for some of its recipients to give the GE team enough time to conduct its baseline survey, which had to be completed before the token transfers were sent. This was a difficult decision, because GiveDirectly does not

want to make recipients wait for their transfers; however, not doing so would have invalidated the randomization for the study, as some community members would have been treated prior to the baseline. The households in Siaya could have been re-randomized, but GiveDirectly had already treated most of them. After debating the pros and cons, GiveDirectly decided to delay the token transfers to preserve the viability of the GE study. This is why the amount of time elapsed between the census and the token transfer in Kenya has recently been greater than usual.

Operational updates

Segovia platform

GiveDirectly expects some efficiency gains from adopting Segovia, but no more than what could be expected from any upgrade in technology, such as an upgrade in finance software. GiveDirectly's chairman, Michael Faye, is the lead on Segovia, while Dr. Niehaus leads on GiveDirectly's work, and each provides occasional support to the other. Combined, the two of them dedicate the equivalent of two CEOs' hours to Segovia and GiveDirectly.

US staffing

To help fill the general management position in its New York office, GiveDirectly has engaged a recruiter, which has supplied several candidates. Some of these are in late-stage conversations with GiveDirectly.

GiveDirectly has seen high interest in its open partnerships lead position, but identifying candidates with the right level of talent has been challenging. Recruiting for this role has been conducted mainly through networking. The candidates GiveDirectly has seen so far would likely take on more of a supporting role, assisting with analysis and strategy. GiveDirectly is also interested in a candidate who could explore more traditionally fundraising-oriented partnerships, such as corporate partnerships that could raise GiveDirectly's retail profile, as well as speak with family foundations and conduct more one-on-one discussions.

Finding candidates for its open digital marketing and products position has also been a challenge, as quality talent in this field is in high demand. GiveDirectly's current management do not have many marketed-related contacts, but have started investigating personal networks more deeply. GiveDirectly has also enlisted another recruiter to help expand the pipeline of candidates.

All GiveWell conversations are available at <http://www.givewell.org/conversations>