

Conversation with GiveDirectly, April 8, 2014

Participants

- Paul Niehaus -- Co-Founder and President, GiveDirectly
- Joy Sun -- Chief Operating Officer, Domestic, GiveDirectly
- Piali Mukhopadhyay -- Chief Operating Officer, International, GiveDirectly
- Carolina Toth -- Kenya Field Director, GiveDirectly
- Eliza Scheffler -- Research Analyst, GiveWell

Note: This set of notes was compiled by GiveWell and gives an overview of the major points made by GiveDirectly staff.

Summary

GiveWell spoke with GiveDirectly about its progress thus far in 2014. Topics included GiveDirectly's current cash disbursement campaigns, operational improvements and experimentation, ongoing and planned research, and GiveDirectly's progress in networking and learning about other cash transfer programs.

This conversation was accompanied by a written attachment prepared by GiveDirectly, referred to throughout these notes as "20140408 GiveDirectly Update," which can be found here:

<http://files.givewell.org/files/DWDA%202009/GiveDirectly/20140408%20GD%20-%20GiveWell%20update.pdf>

Table of contents

- I. Summary
- II. Revenue and expenses in 2014
- III. Cash disbursement
- IV. Village selection mechanisms
- V. Operational improvements
 - a. Communication at enrollment
 - b. Preempting adverse events
 - c. Payments
 - d. Staff management
- VI. Experimentation
 - a. Saturation
 - b. Community-based targeting
- VII. Research
 - a. Planned research
 - b. Topics for potential future studies
 - c. Past research
- VIII. Networking

Revenue and expenses in 2014

GiveDirectly plans to spend \$11.6 million in 2014 on field (\$10.8 million) and domestic (\$0.8 million) expenses. This is lower than the \$13 million target that GiveDirectly set at the end of 2013; the reason for reducing its target spending for 2014 is to leave some funds on reserve to buffer for uncertainty regarding fundraising in the future, so that GiveDirectly will be able to retain its core staff even if future revenue falls short of expectations. GiveDirectly has raised a total of \$15.1 million so far this year, including \$10 million from the Good Ventures match. This revenue fully funds GiveDirectly's planned expenses for the year, with some left over for reserves.

Cash disbursement

GiveDirectly's planned disbursement of cash transfers can be divided into three campaigns (more information on "20140408 GiveDirectly Update," Pg 3):

1. **Kenya 1.2M campaign:** 1,200 recipients, enrollment initiated November 2013, projected completion date November 2014; total budget \$1.4 million
2. **Kenya scale-up:** 6,500 recipients, enrollment initiated March 2014, projected completion date July 2015; total budget \$7.3 million
3. **Uganda scale-up:** 2,000 recipients, enrollment initiated April 2014, projected completion date May 2015; total budget \$2.3 million

Village selection mechanisms

In each of the three campaigns, villages have been selected for enrollment using different mechanisms. In the Kenya 1.2M campaign, villages were selected based on the proportion of thatch to iron roofs, which was manually estimated using satellite imagery. In the Kenya scale-up, GiveDirectly used a machine learning algorithm that estimates the proportion of thatch to iron roofs in a village using satellite imagery. The algorithm was developed for GiveDirectly by a group called DataKind, which provided its support pro bono.

In the Uganda scale-up, GiveDirectly selected villages based on census data at the parish level (a parish has a few villages). The data allowed GiveDirectly to identify very poor districts in Eastern Uganda, its target region. This data is not easily publicly available; GiveDirectly gained access to it with the help of a group of data scientists at Stanford who are providing pro bono support to GiveDirectly to improve its targeting methodologies.

Operational improvements

The following section expands on “20140408 GiveDirectly Update,” page 4.

Communication at enrollment

Feedback from qualitative surveys of Nike recipients and conversations with village members in other campaigns demonstrated that some recipients felt they would like more support in making important decision about how to spend their transfers. There were also some recipients who believed that transfers were meant to be spent in relation to housing, due to the selection criteria, despite existing messaging that transfers are unconditional. To address this, GiveDirectly decided to augment its communications with recipients at enrollment to include more information about a broad variety of spending options, provide some planning tips, and emphasize that the transfers are unconditional.

GiveDirectly added questions to the script that field staff use at enrollment such as “have you thought about how you will use the transfer?” and “have you spoken with your spouse about how you will use the transfer?” Starting with the Kenya \$1.2M campaign, recipients now receive a document showing different categories, written in the local language, of items that they can spend funds on, with images accompanying each category (see attachment “Registration menu”). GiveDirectly noted the possibility that part of the reason the vast majority of recipients have spent funds on housing materials is that housing materials are the basis for the household targeting criteria. The purpose of this menu is to help recipients think about different options for spending, while field staff still emphasize that households can spend the money however they’d like. Field staff who worked on the Kenya \$1.2M campaign reported that many recipients valued having the menu. GiveDirectly plans to update the menu with more realistic images (photographs of actual items) and may include prices in the future.

Preempting adverse events

GiveDirectly would like to better understand the cases of conflict that have arisen between recipients and their family or community members, in order to inform its process for conflict resolution and mitigation. To do this, one of GiveDirectly’s senior staff or field officers who has strong communication skills is going to conduct conversations with people who have been involved in these adverse events.

Another change that GiveDirectly has implemented in Kenya to preempt adverse events is identifying people who may be especially vulnerable to having problems with receiving and spending their transfers. During the back check, which is GiveDirectly’s

last in-person step before transfers are sent, field staff nominate people who they consider to be vulnerable cases. This may include people who appear mentally unwell, disabled, very old, or anyone who requires a trustee to receive the transfers on their behalf because the government will not issue them formal identification. (Trustees are selected by the recipient or the recipient's family if the recipient is not able to do so).

After the initial token transfers are sent, GiveDirectly follows up in person with all of the people identified as potentially vulnerable. If their cases continue to seem worrisome, GiveDirectly field staff will follow up in person again after the first lump sum is sent. The first round of follow up has been completed for the Kenya \$1.2M campaign, primarily by one field staff who is a former counselor and one who formerly worked for the Centers for Disease Control and Prevention (CDC). Out of the 1200 total households in that campaign, there were 10 cases of vulnerable recipients having problems that GiveDirectly was able to help with and believes it would not otherwise have found. An example of one of the problems faced by a vulnerable recipient is that his trustee, who was his brother living in the village, was not trustworthy. When GiveDirectly spoke with the recipient, he requested that his other brother, who lived in a nearby town, be made the trustee; GiveDirectly called this brother, and he came to the village to be registered as the official trustee.

Payments

Flexible transfer timing in Kenya

In Kenya, GiveDirectly would like to offer flexible timing of transfer payments as an option for recipients, and plans to evaluate the impacts of this as part of the behavioral economics study (see "Research" section). GiveDirectly is working on its transfer automation technology to make sure it can accommodate flexible timing, and has also written a script for field staff to reduce the inevitable confusion when people are receiving transfers at different times, and when some people have a choice about timing while others do not. GiveDirectly is also increasing its capacity to manage a higher volume of calls to its hotline.

Implementing lump sum payments in Uganda

(This is a possible future operational improvement, though it is not currently in GiveDirectly's workplan.)

In Uganda, transfers are sent in ten equal installments, and GiveDirectly works with the mobile money providers to coordinate pay out days, as their agent networks are less robust than those in Kenya. GiveDirectly expects to continue to have to coordinate pay

out days, but hopes that it can move towards sending larger payments less frequently, like the Kenya model of a token transfer then two lump sums. Past recipients have expressed a preference for lump sums because this enables them to more easily make large investments. In addition, it reduces the time that GiveDirectly has to spend on coordinating and overseeing the pay out days, though this has become more streamlined over time. GiveDirectly continues to seek other mobile money providers who it might want to partner with for future campaigns in Uganda (e.g., Airtel).

Biometric authentication

(This is a possible future operational improvement, though it is not currently in GiveDirectly's workplan.)

The government of Uganda started a large cash transfer program called Social Assistance Grants for Empowerment (SAGE), which provides \$20 monthly transfers to eligible people in Uganda. The program currently serves 100,000 households in 17 districts and has plans to scale up; it is not currently active in Bukudea, where GiveDirectly operates. The government of Uganda is working with the mobile money provider MTN to build the capability to use biometric authentication (fingerprinting) for transactions and account access. GiveDirectly is interested in running a pilot of biometric authentication with its own cash transfer recipients who are serviced by MTN.

Staff management

Additional management staff

GiveDirectly recently hired Lawrence Juma to serve in a new role it has added to its Kenya staff: Project Associate. This role is responsible for directly managing and coaching field staff and overseeing data collection, which will alleviate these responsibilities from the Kenya Field Director. Mr. Juma was previously involved in project management at Innovations for Poverty Action. The current Field Director in Kenya, Carolina Toth, will soon be transitioning to a domestic role on GiveDirectly's staff, and the current Deputy Field Director, Joe Huston, will be taking her place.

GiveDirectly has hired multiple staff members to fill key roles, including full management capacity in Uganda, who would be able to shift to Kenya in the event that GiveDirectly decides to allocate a higher proportion of transfers in Kenya in the future. GiveDirectly is still looking to hire more technology staff.

Incentivizing staff performance

GiveDirectly developed a system to evaluate field staff performance, using data on productivity and quality of work combined with subjective measures of teamwork and cooperation, as assessed by field staff supervisors. The system will generate monthly “scores” for each field staff member. GiveDirectly is not going to reduce the baseline salaries that it offers to staff, but it will offer bonuses to those with the highest performance scores. The evaluation system has been well received in concept by GiveDirectly field staff; it is about to be implemented for the first time in Kenya.

Eliminating time lags between enrollment periods

In previous campaigns, GiveDirectly completed each step of the enrollment process in sequence, and enrollment teams did not work at the same time. Due to the much larger size of the current campaign, GiveDirectly has shifted to a rolling model in which all three steps occur simultaneously across a much larger range of villages. Field staff now work continuously, year-round as long as funds are available to distribute, rather than in discreet campaigns as before.

A major benefit of the rolling model is that it will increase the number of recipients that can be registered in a given time period. It will also enable GiveDirectly to employ field staff full time and have them relocate for the job, rather than offering temporary employment and paying for travel and accommodations during the working period. Full time positions make it easier for GiveDirectly to retain its best staff members.

However, the rolling model also presents some new challenges. For one, more recipients registered means more data to manage. Also, GiveDirectly must calibrate the field team sizes so that they can complete their work on the same schedule. There is also an increased risk of collusion among staff members who will now be living in the same city and will know each other better. As GiveDirectly scales up, it’s management staff will not know the field staff as well, so it needs to design a system that prevents against schemes such as field staff teams agreeing to enroll an ineligible household in exchange for part of the transfer funds. One regulation that GiveDirectly has put in place is strong restrictions around communication of recipient information between field teams. It is also planning to make sure that field teams are never working in the same place on the same day.

Joe Huston, the incoming Kenya Field Director, is working to automate much of the data processing that happens during enrollment, such as integrating data collected at different times, generating lists of households and assignments for staff. GiveDirectly hopes that automation will reduce the time spent and errors made in these processes.

Experimentation

The following section expands on “20140408 GiveDirectly Update,” page 5.

Saturation

In its standard model, GiveDirectly provides cash transfers only to the households that have thatch roofs. GiveDirectly experimented with more inclusive targeting in 19 randomly selected villages, in which nearly all households received transfers (all except those made from fully permanent materials such as cement walls and iron roofs). GiveDirectly compared these villages to 18 villages in the same region where standard targeting was applied. The factors being compared were cases of conflict/tension reported in follow-up surveys and focus groups, and instances of gaming that were discovered by GiveDirectly field staff throughout the cash transfer process.

GiveDirectly found that conflict and tension were not significantly lower in the saturation villages. The most common form of tension mentioned was social awkwardness between recipients and non-recipients, such as avoiding eye contact when coming across each other. This type of tension was mentioned in both standard targeting villages as well as saturation villages, where it occurred between members of the recipient village and members of neighboring villages when crossing paths in market centers. People in both villages predicted that the social awkwardness would pass once the money was spent.

GiveDirectly found low levels of gaming across villages (fewer than one instance per village) and did not detect a difference between standard and saturation villages. Of all of the households initially marked as eligible during the census for the 37 villages, 4% were determined ineligible at a later point in the enrollment process.

When GiveDirectly asked village members if it was better to give transfers to all the members in one village, or to only thatch-roof households in two villages, everyone picked the latter. Village members said that targeting thatch-roof households was a good way of identifying the poorest. Moving forward, GiveDirectly is planning to continue applying its standard targeting criteria. However, it may revisit this idea if it starts working in villages with a very high proportion of thatch-roof to iron-roof households (>80%), because these villages may be so poor that no households should be excluded. The villages in which GiveDirectly currently operates have about 40% thatch-roof households.

GiveDirectly shared a summary of the results of its experimentation with saturation targeting (see attachment "Saturation analysis").

Community-based targeting

GiveDirectly conducted village meetings in two villages before recipients were enrolled, to determine community preferences on eligibility criteria. In both of these pilots, village members chose housing materials as the best indicator of poverty. In a follow up meeting, GiveDirectly asked village members to work in groups to categorize each household in their village as having a thatched roof, iron roof, or made from fully permanent materials. About 6 groups of 10-15 village members each participated in the household classification; each group was facilitated by a GiveDirectly Field Officer. The lists of households provided to the groups for classification were prepared by a few village members, including the village elder and a community health worker, under the advisement of one of GiveDirectly's Senior Field Officers.

GiveDirectly used the categorized lists generated from these village meetings in place of a census. Households that had been marked as eligible by all groups were added to the registration list, households unanimously marked as ineligible were not added to the registration list, and in cases where groups had made different categorizations of the same household, GiveDirectly staff followed up with the household to determine eligibility. GiveDirectly's impression is that the village members who participated in community-based targeting enjoyed the process. Another benefit is that it could be less expensive than a census for GiveDirectly to administer the household classification at a village meeting.

The next iteration of community-based targeting that GiveDirectly would like to try is adding a step at the end of the enrollment process where village members can nominate recipients to receive transfers who they believe were wrongfully excluded from the census or are especially deserving despite living in iron-roofed households (e.g., a disabled widow). GiveDirectly expects this step to increase the enrolled households by no more than 10% and feels it is important to provide an opportunity for community input in its targeting process. At this point, GiveDirectly is planning to implement this additional step in 1-2 villages. GiveDirectly is learning from what other NGOs have done with similar activities and is also getting feedback on its process in order to make the village meetings as participatory and fair as possible.

Research

GiveDirectly is currently moving forward with three formal impact-oriented research projects, which are briefly described on page 6 of "20140408 GiveDirectly Update." The budgets noted in that document include the costs of cash transfers and implementation of the transfer campaigns, which are paid for by GiveDirectly, as well as the costs of

research (staff time to conduct measurement and analysis), which are not paid for by GiveDirectly.

Planned research

Study of behavioral interventions

This study will assess the impact of transfer timing and information about social norms on the process by which recipients' decide how to spend transfers, long-term outcomes for recipients such as income and assets, and recipients' aspirations for the future. The size of this study is contingent on the amount of funding secured. An anonymous funder has agreed to provide a grant for the study, which may be used to cover both implementation and research costs of a smaller study involving approximately 500 recipients. If additional funding is secured to cover implementation costs, the grant from the anonymous funder will be applied only to research costs, and the study would be conducted at a larger scale of approximately 2000 recipients. If the study is initially conducted at the smaller scale, it could still be scaled up over time as funding becomes available.

The research for this study will be conducted by Anuj Shah and Sendhil Mullainathan of Ideas42. The researchers have a deep interest in psychology and poverty, and are committed to exploring many questions on this topic as well as on evaluation design. A behavioral study of GiveDirectly's program could help prioritize future research and attract funders for this work.

Study of general equilibrium effects

This study, which will assess the impact of large-scale cash transfers on local economies, is being led by Professor Ted Miguel and a graduate student at the University of California, Berkeley. This team has conducted power calculations to determine what size sample will be necessary to detect effects on the local economy. If the sample is randomized at the village level, the study would need to include about 15,000 households (\$15 million in transfers). If it is randomized at the sub-location level, the required number of households would increase, as sub-locations include a few villages each.

A proposal for \$5.5 million of funding has been submitted to a potential funder. Of this, about \$200,00 - 300,000 would be used for research costs, which would scale linearly as the study is expanded. GiveDirectly has already been randomizing its selection process of villages eligible to receive transfers in Kenya so that these villages could be incorporated into this study. GiveDirectly plans to continue to randomize villages that

receive transfers, but noted that it would be good if the researchers could take baseline measurements in these villages before they receive transfers.

Study of intra-household bargaining dynamics

This study will test ways of altering the cash transfer disbursement process to improve the dynamics of conversations within households regarding how to spend the funds. GiveDirectly said that the Rarieda RCT demonstrated that the cash transfers had positive effects on reducing domestic violence, violence against women, and conflict within the household, however there have also been isolated cases of conflict between spouses in GiveDirectly recipient households. This study will test different approaches, such as facilitating planning conversations between spouses before transfers are sent and having couples sign agreements. Researchers familiar with this field of work have suggested that these strategies could positively influence intra-household bargaining dynamics.

This study is not yet funded, but GiveDirectly is in conversation about it with one large funder.

Topics for potential future studies

There are other questions that GiveDirectly deems important and would like to explore through formal evaluation, though nothing has been planned for these studies at this time.

Topics:

- *Long-run household impacts*: extended data collection beyond the time frame of the Rarieda RCT
- *Transfer size and lumpiness*: some relevant data was generated from the Rarieda RCT, but those findings are merely suggestive, as the study was powered to detect an average treatment effect and not sufficiently powered within treatment variations to demonstrate differential effects with specificity. GiveDirectly would like to revisit this at a larger scale.
- *Heterogeneity of impacts for differing recipient populations*: such studies could assess the returns by recipient income segment (i.e., more/less poor); impacts for girls/young women (building on Nike pilot); or health impacts of cash transfers (e.g., targeted at women with children under 5). GiveDirectly is planning to use the data from the Rarieda RCT to see what can be learned about impacts between different types of recipients. However, given that this study included only 500 recipients, GiveDirectly does not expect that it will be able to draw

strong conclusions. There is also a question of whether cash transfer programs *should* be targeted to reflect these findings.

Past studies

Extended data collection by mobile phone

IPA conducted a study of collecting follow up data from GiveDirectly recipients by mobile phone, which was funded by USAID via the Policy Design and Evaluation Lab at UCSD. The main findings from this study were that collecting data via voice calls works well, yielding high compliance and clean data, while collecting data via SMS was not as successful, resulting in more attrition and less useful data.

Networking

GiveDirectly is putting a lot of effort into building its network in Kenya and Uganda, including public sector officials who administer government cash transfer programs, development aid agencies and other large aid organizations that fund cash transfers (e.g., the UK's Department for International Development, USAID, UNICEF, the World Bank), and with non-profit organizations that work on cash transfer programs (e.g., Oxfam International, CARE).

GiveDirectly's board members and International Chief Operating Officer have been meeting with representatives of these groups to learn about the implementation challenges they face and how they've addressed them. These meetings are also a way for GiveDirectly to disseminate what it has learned about operations and its research findings on the impacts of GiveDirectly's transfers. Networking also enables GiveDirectly to explore whether there is potential for it to work in partnership with other groups, such as providing technical assistance for program design or serving as the implementer for part of a government cash transfer program.

GiveDirectly is also networking with people who are active in the civic space in Kenya and Uganda - thought leaders, journalists, and activists. GiveDirectly wants to share information with these people because they influence the public dialogue. GiveDirectly is also meeting with senior or retired government officials who can provide guidance on navigating the government and connect GiveDirectly to allies on the public sector side.

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