

Researchers

Joshua Blumenstock
University of California, Berkeley

Dean Karlan
Northwestern University

Christopher Udry
Northwestern University

Research Implemented by IPA

Yes

Using Mobile Phone and Satellite Data to Target Emergency Cash Transfers in Togo

Abstract

Since March 2020, a team of researchers¹ has been providing technical assistance to the Government of Togo to help guide their humanitarian response to the COVID-19 crisis. This page provides a short, non-technical summary of those efforts.

Background: Togo and the COVID-19 Crisis

In Togo, a small country in West Africa, over 50% of the population lives in poverty. The COVID-19 pandemic threatens to reverse years of poverty reductions, as lockdown measures designed to stop the spread of the virus limit economic activity and threaten food security. More broadly, the World Bank [estimates](#) that the pandemic will force between 88 and 115 million people in low and middle-income countries into extreme poverty in 2020.

In response to this crisis, the Government of Togo launched “[Novissi](#),” an innovative social assistance program designed to provide emergency cash assistance to Togo’s neediest families. Over the course of just a few weeks, the government built and deployed a completely contactless, digital system that provided over half a million individuals with cash. Beneficiaries registered using their mobile phones (no smartphone was required); after entering basic information into a USSD menu, they were then immediately sent mobile money transfers of approximately \$15/month, lasting for three months. This program was the first of its kind, [described as](#) an “exemplary case of social protection in response to the coronavirus pandemic in Africa.”

The Targeting Challenge

A key question at the heart of Novissi—and at the heart of the hundreds of other targeted social protection programs that have been launched in response to COVID-19—is *how to prioritize those people with the greatest need?*² The difficulty was that the Togolese Government did not have a comprehensive social registry that would allow them to directly identify and prioritize its poorest people. The last census was conducted in 2011, and even that database did not have any information that could allow the government to determine who should be prioritized for assistance. And in the middle of a pandemic, it was impractical to collect the data required to create a new registry.

When Novissi first launched in April 2020, the government chose to prioritize informal workers living in the regions most impacted by the lockdown restrictions. To determine eligibility, they relied on a voter database that had been recently updated, and in which people had self-declared their home location and their occupation.

Guiding the Expansion of Novissi

The Government of Togo, in partnership with GiveDirectly, is currently expanding Novissi to rural areas of the country, where extreme poverty is most severe. A top priority in this expansion is ensuring that benefits are targeted toward the poorest citizens. A team of researchers from the Center for Effective Global Action and Innovations for Poverty Action are supporting this effort. The novel approach developed by the research team has two main stages.

Stage 1: Prioritizing the Poorest Villages and Neighborhoods

The Government of Togo was interested in prioritizing support to the poorest cantons (admin-3 regions) in Togo. However, prior to the collaboration with the research team, the Government of Togo did not possess a poverty map that would allow for such granular targeting. The government had previously conducted nationally representative household surveys, but those surveys could only produce estimates of poverty representative at the national (admin-0) or regional level (admin-1). To enable more granular geographic targeting for the program, the research team used the deep learning pipeline described in Chi et al. (2020) to identify the poorest 100 cantons in the country (the number 100 was chosen by the government partners, based on the estimated distribution of wealth of the population living in those cantons).

Specifically, the research team started with a large, nationally-representative household survey that had reliable estimates of the living conditions of thousands of households. The research team used those survey-based estimates of poverty as “ground truth” to train a machine-learning algorithm to estimate the wealth of very small regions (i.e., 2.4km tiles) based on the geographic characteristics of that region—where the researchers used rich data from satellites and other sources to measure those geographic characteristics. Intuitively, the algorithm learns that certain patterns in the imagery are indicative of wealth (such as places

with metal roofs and high-quality roads) while others are indicative of poverty (such as places with certain types of terrain and weather). The research team then overlaid tile-level consumption estimates with high-resolution population density estimates from [Tiecke et al. \(2017\)](#) to estimate average per capita household consumption for each of Togo's 397 cantons. The tile-level and canton-level estimates are shown in Figure 1 below. These maps were used to determine the 100 poorest cantons.

Figure 1: Prioritizing the Poorest Villages and Neighborhoods



Figure Notes: The research team produced micro-estimates of the wealth of each 2.4km grid cell by applying deep learning algorithms to high-resolution satellite imagery (left), combined those estimates with information on the population density of each grid cell (middle), and used this information to determine the 100 poorest cantons in Togo (right).

Stage 2: Prioritizing the Poorest Individuals in the Poorest Villages

In expanding Novissi to rural areas, one of the government's goals was to ensure that benefits flowed to the poorest individuals living in the 100 poorest cantons. Specifically, in collaboration with GiveDirectly, the government secured sufficient funding to provide benefits to roughly 57,000 of the approximately 580,000 citizens living in the poorest 100 cantons. Building on recent advances in the research literature, the research team developed methods that use mobile phone metadata to help identify the individuals with the greatest need.³

Specifically, with financial support from the World Bank, the research team conducted a large phone survey in September 2020, immediately prior to the expansion of Novissi in rural areas, to provide "ground truth" information on the living conditions of roughly 10,000 individuals. The research team went to great lengths to ensure that the individuals in this phone survey were representative of the broader population of people living in the 100 poorest cantons, through a combination of adaptive survey design and the use of survey sample weights. This helps ensure that difficult-to-reach populations, such as the extreme poor and those living in remote villages, are represented in the training data. Each survey respondent provided informed consent to participate in this research program.

For each of the 10,000 survey respondents, the research team then matched information they provided on their daily consumption to their mobile phone metadata, which was obtained from the mobile phone operators in Togo—see below for a discussion of how the researchers protected data privacy. The metadata includes information about the date, time, duration, and cell tower used for calls and texts, as well as information about each subscriber's volume of mobile data usage and mobile money transactions. The research team processed these raw data with algorithms that construct aggregate statistics of each subscriber's pattern of mobile phone use, including information that is correlated with wealth, such as the total volume of international phone calls made, or the average mobile money balance. On an anonymized version of this dataset, the research team then trained machine learning algorithms to predict consumption from mobile phone metadata. These algorithms were then used to generate a consumption estimate for each of the 5.7 million mobile

subscribers in the country.

Figure 2: Prioritizing the Poorest Mobile Subscribers



Figure Notes: Using ground truth wealth and poverty collected through a large phone survey of active mobile phone subscribers, the research team trained machine learning algorithms to estimate the wealth of each mobile subscriber (left). In the 100 poorest cantons (red distribution in right figure), those estimated to consume less than \$1.25/day are prioritized for Novissi (dashed vertical line). These individuals are substantially poorer than the average resident of Togo (blue distribution).

Preliminary Findings

After developing this approach, the research team carefully and critically evaluated it using several independent datasets. The team focused in particular on comparing this new targeting approach to alternative approaches that could have been used by the Government of Togo at the time. In particular, for the expansion of Novissi to rural areas, the Togolese government was considering two alternatives in addition to the phone-based approach: an expansion of the existing occupation-based targeting; and a “geographic blanketing” approach that would provide benefits to all individuals residing in specific villages.

The research team used data from the September 2020 phone survey to compare these options and evaluate the effectiveness of each of the different methods in reaching the poor. Preliminary results indicate that assuming the goal is to reach the poorest 57,000 people in the 100 poorest cantons, the satellite+phone approach is significantly more accurate than the alternative approaches available to the government. In particular, as shown in the figure below, the satellite+phone approach is expected to provide benefits to nearly 2.5X as many of the poorest citizens as a program that provided benefits based on occupation.

Figure 3: Comparison of the Satellite+Phone Approach to Existing Alternatives



Figure Notes: In simulations, the research team assumed each of the three possible approaches to targeting benefits has a budget constraint of paying no more than 57,000 beneficiaries. Figure shows the accuracy of each mechanism at reaching the poorest (i.e., the lowest consumption individuals in the 100 poorest cantons), simulated using a phone survey of 9,484 individuals living in those cantons, conducted in September 2020.

Ongoing and Future Work

Through a collaboration with the Government of Togo and GiveDirectly, roughly 57,000 new beneficiaries are being enrolled in Novissi using the satellite+phone method described

above. The research team is currently conducting a careful evaluation of this program expansion. In particular, they have planned in-person surveys for Spring 2021—or as soon as public health restrictions allow—to talk to program beneficiaries, community members, and village leaders. These interviews will make it possible to critically assess the program and compare the new targeting methodology to common alternatives used in the design of anti-poverty programs, such as a proxy means test (PMT) and community-based targeting (CBT). Although the Togolese government was not in a position to use a PMT or CBT for the Novissi program, they represent important benchmarks for comparison.

Key Concerns and Limitations

The research team is optimistic that this new approach to targeting can provide an effective means to distribute cash to extremely poor people, extremely quickly. At the same time, there are important practical, social, and ethical considerations involved in the real-world implementation of such a program—these considerations were a major focus of the researchers' time and energy as they worked with the Government of Togo and GiveDirectly to design a humanitarian aid program based on this technology. Here the team highlights a few key concerns, as well as their approach to dealing with those concerns.

Program exclusions: Of central importance to any targeting mechanism, including the one developed by the research team, are exclusion errors (people who need benefits but don't receive them) and inclusion errors (people who don't need benefits but do receive them). There is a large academic literature that documents the targeting errors that are likely to occur with existing targeting methods (cf. [Banerjee et al. 2009](#), [Alatas et al. 2012](#), [Alatas et al. 2016](#)). Thus, a major focus of the in-person survey the research team has planned for early 2021 is to rigorously document the targeting errors that occurred in this initial roll-out—to make sure such errors are transparent and can be addressed in the future.

For this program, one important concern is exclusions that occur for people who don't have mobile phones. While this type of exclusion is not a product of the researchers' new method—it exists by virtue of the fact that the only way the government could quickly distribute cash *en masse* during the pandemic was by using mobile money—it is a concern that the research team is tracking very carefully through surveys, audits, and engagement with community leaders (see below for more on this). Based on the research team's analysis of recent nationally-representative household survey data, the team estimates that roughly 90% of households in Togo have at least one mobile phone, which might limit the scope of such exclusions. In addition, the current expansion was designed to allow anyone with a *SIM card* to register for Novissi. In other words, someone without a phone could acquire a SIM card (which is much cheaper than a phone) and use that SIM card in someone else's phone to register for the program and receive benefits. In future expansions, the research team is exploring the possibility of making low-cost mobile phones available through [pay-as-you-go](#) subsidies.

The research team has also taken several measures to advertise the program to the public and encourage enrolment. For instance, the team has briefed village leaders about the program, sent hundreds of thousands of SMS messages to the intended population of

beneficiaries, distributed posters to hundreds of communities, and placed advertisements on local radio stations. The researchers' hope is that these efforts to promote program awareness can help reduce the scope for exclusion.

Data privacy: In designing this program, user privacy was a top priority. The research team took several steps to minimize the use of private data to the greatest extent possible and had UC Berkeley's Committee for the Protection of Human Subjects review all of the research procedures involved in this work. Among other protections, the research team built in organizational safeguards that minimized the data to which each party had access: for instance, neither GiveDirectly nor the Government of Togo ever has access to any data collected by the mobile phone operators. Likewise, neither the government nor GiveDirectly received access to the poverty scores derived from the mobile data. Instead, the research team produced a list of beneficiaries based on the poverty scores, and those beneficiaries were made eligible to receive transfers. The only data the government received was this list of eligible beneficiaries. The researchers also implemented strict anonymization, encryption, and access protocols, for instance to ensure that all data are pseudonymized to remove personally identifying information prior to analysis. Third, they collected consent whenever possible: for instance, phone survey respondents provided informed consent before participating, and program beneficiaries are asked for consent to use their data for program administration during the USSD registration process.

Fairness and bias: The research team recognizes the fact that all algorithms are likely to be biased. As noted earlier, the team was careful to design the survey and sampling strategy to ensure that the data used to train the machine learning algorithm were as representative as possible of the broader population of mobile phone subscribers in the 100 poorest cantons. The team also used several independent sources of survey data to calibrate and double-check the output of these algorithms, and to conduct algorithmic audits to examine if specific vulnerable subgroups were more likely to be excluded on the basis of the algorithm. The research team has also planned extensive survey activities for early 2021 that are designed to determine if the algorithm's predictions were inadvertently biased against any particular type of beneficiary (such as women, illiterate, and marginalized subgroups); these surveys will also help determine if this method is better at reaching some populations that are often excluded through alternative approaches to targeting the poor (e.g., community-based targeting can often neglect socially excluded individuals). Still, the researchers are aware that not all types of bias can be anticipated, and thus are actively engaging with community members through qualitative interviews and other feedback mechanisms to try and surface any unintended biases or exclusions created by this program.

Community outreach and support: Prior to the launch of the Novissi in rural areas, and after extensive consultation with partners in Togo, GiveDirectly and members of the research team visited several communities to conduct qualitative focus groups and interviews with likely beneficiaries. The research team also worked with the government to contact more than 100 village and community leaders, to let them know about the program, and to create a relationship to receive feedback on program design and implementation.

As noted, the research team also took several measures to communicate the program to the

public, through segments that were aired on local radio, as well as broadcast SMS messages and public posters. The team established a call center that provides 24-hour real-time support and troubleshooting, and beneficiaries can also directly submit feedback via the phone-based USSD platform. The researchers also have several real-time auditing mechanisms: the team is constantly monitoring program data to ensure there are no issues with registration or payments; the government has hired Ernst & Young as external auditors; and GiveDirectly continually conducts qualitative follow-up calls with program participants and community members to monitor program integrity.

Summary

In partnership with the Government of Togo and GiveDirectly, researchers have helped build an innovative approach to humanitarian crisis response, designed to help as many people as possible, as quickly as possible. By leveraging new sources of digital data and recent advances in machine learning, they demonstrate how to target humanitarian cash transfer relief programs to people in need, even in settings where no universal social registry exists.

Acknowledgments

The research team is grateful to their project partners in Togo who have supported this initiative, especially Minister Cina Lawson, Shegun Bakari, Stanislas Telou, Leslie Mills, Kafui Ekouhoho, and Attia Byll. The research team at GiveDirectly has also been instrumental in supporting this effort, especially Han Sheng Chia, Michael Cooke, and Alex Nawar. The research team is supported by invaluable research assistance from Isabel Falomir, Shikhar Mehra, Suraj Nair, Nathaniel Ver Steeg, and Rachel Warren. Tina George and Zoe Kahn provided helpful input on this write-up.



Financial Support



¹ Research team: Emily Aiken, Suzanne Bellue, [Joshua Blumenstock](#), [Dean Karlan](#), and [Chris Udry](#). For questions, [contact Joshua Blumenstock](#).

² This question of *program targeting* is one of the central difficulties in the design of anti-poverty programs ([Coady et al. \(2004\)](#), [Brown et al. \(2018\)](#), [Hanna and Olken \(2018\)](#)).

³ This approach is based on methods described in greater detail in [Blumenstock et al. \(2015\)](#) and [Aiken et al. \(2020\)](#).

January 12, 2021