Muralidharan, Niehaus, Sukhtankar and Weaver
June 1, 2018

Pre-Analysis Plan for High-Frequency Program Monitoring and Bureaucratic Performance: Experimental Evidence from India[1]

A. Introduction

Throughout the developing world, citizens often struggle or fail to acquire public services to which they are legally entitled. Unlike those who deliver services in the private sector, government bureaucrats delivering public services lack a profit motive to improve customer experience. As a result, there is a greater need to build in systems that discipline public servants in charge of delivering public services. Yet government monitoring efforts in developing countries are typically constrained by slow, indirect monitoring processes and unrepresentative, sporadic data. In this study, we evaluate the effects of a method for collecting high-quality, real-time data on program implementation: a high-frequency monitoring (HFM) system based on outbound phone calls to a large, representative sample of program beneficiaries.

Under the HFM system, a call center places phone calls to program beneficiaries that collect information on their experience with specific government programs. This information is collated and provided to relevant parties. In the context of our evaluation, this type of information aggregation system can have two main effects:

- **Incentive effect:** Motivate better performance from frontline officials by providing reliable information on their performance relative to others. Typically, even if more senior officials are interested in monitoring and sanctioning poorly performing frontline workers, they lack the data to do this successfully.
- **Information effect:** Generate real-time data on local-specific problems for resolution by local officials.

We study the effects of such a system in the context of government cash transfer program in the Indian state of Telangana in May 2018. This program, known as the Rythu Bandhu Scheme (RBS), is a flagship initiative of the state government that provides land-owning farmers with Rs.

[1] This pre-analysis plan registers our analysis on data from two sources: high frequency monitoring phone calls and government records. We will later file a pre-analysis plan for future rounds of data collection. Since the high frequency monitoring phone call data and government records data will soon be available to the researchers, it is necessary to file this now.

4000 (approximately USD60) for each acre of land that they own[2]. This money is intended to be distributed prior to each growing season and used for agricultural inputs such as fertilizer and seeds.

In order to implement RBS, the government of Telangana (GoTS) updated all of the land records for the state with the current owners' name and a government unique identification number (Aadhaar). Money is distributed to farmers in the form of "order cheques", which can be exchanged for cash at any bank branch of the bank listed on the cheque. To receive the cash, the farmer must show an official form of identification that matches the name and unique id number listed on the cheque.

The distribution of these cheques was supposed to be carried out at a series of village meetings over the course of 12 days (May 8 to May 20, 2018). The individuals distributing the cheques are agricultural extension workers, who are overseen by "mandal agricultural officers". These mandal agricultural officers are the focus of our intervention and oversee the Department of Agriculture work within the level of a "mandal" (a geographical agglomeration of approximately 65,000 people).

Given the value of money being distributed through the scheme (approximately US$1.77 billion per year), a natural concern is whether it reaches the intended beneficiaries. Possible implementation issues are that cheques do not reach the intended beneficiaries, that there may be corruption during the distribution process (e.g. those distributing the cheques demand payment for handing them out), or that there would be delays in distribution, a major concern given the need to purchase agricultural inputs within a particular time window.

Mandal agricultural officers (MAOs) were randomly assigned to either a treatment or control group as detailed in the randomization section. We compare across the groups to determine the effect of implementing a high frequency monitoring system.

In the treatment group:

- Prior to the distribution of cheques and start of the HFM system, the Telangana Department of Agriculture held a video conference with treatment MAOs to inform them that they were selected for a pilot initiative. During this meeting, the state Commissioner of Agriculture explained the HFM system and the types of data that would be collected. After the meeting, the Department of Agriculture sent a letter to all of the treatment MAOs with the same set of information in written form.
- The video conference and letters also informed the MAOs that multiple reports from the HFM system would be provided to them and their superiors, including an implementation performance rating for their mandal relative to other nearby mandals. They were not told exactly how the rating would be constructed, aside from the types of data that would go into it.

[2] This was limited to farmers with less than 100 acres. In practice, out of a total population of 5.7 million registered landowning farmers, there were only 114 farmers who owned 100 or more acres of land.

- In order to reduce the risk of spillovers, treatment MAOs were explicitly told the identity of other treatment MAOs in their district and that no other mandals in their district are part of the pilot program.
- After the HFM data was collected, treatment MAOs and their superiors were sent a report on program implementation in their mandal and a rating of this mandal relative to other treatment mandals. This report will have information on some villages where implementation was particularly poor.
- After the initial report was issued, the HFM system collected a second round of performance data two weeks later. The MAOs and their superiors received reports on whether the implementation situation has improved in those mandals.

In the control group:

- These MAOs were not be explicitly informed about the existence of the HFM pilots in other mandals. If they asked, they were told that the data collection may occur in future rounds of RBS in their mandals, such as the next season, but not in the present round.
- HFM data will still be collected from farmers in control mandals, but this data will not be used to generate reports on the performance of control MAOs, and control MAOs will not be made aware of this.

The high frequency monitoring call center collects information on: 1) whether the farmer received their cheque; 2) date of cheque receipt; 3) whether they encashed their cheque; 4) problems in receipt of the check (e.g. asked to pay money, time cost); 5) problems in encashing of cheque; 6) satisfaction with the scheme; 7) other feedback on the scheme. Calls are made via a combination of human phone surveyors and interactive voice response calls, where recipients indicated their responses via a touch-tone system. This information is compiled into reports to state, district, and mandal level officials. The reports do not identify individual phone respondents for the sake of their privacy, only aggregate information.

This document describes our plan for measuring the impacts of the HFM system, including the design for randomization and for sampling, primary and secondary outcomes of interest and our plan for measuring them, other descriptive measures we will calculate, and statistical methods we will use to analyze each measure.

B. List of Data Sources

Our primary outcomes of interest relate to whether the treatment improved delivery of cheques. The data will be collected in three rounds, with each providing a different set of insights.

The first round of data is composed of the first round of high frequency monitoring data and government monitoring and information system (MIS) records as of the end of the first round of HFM. The treatment is the revelation of the high frequency monitoring data and informing the MAOs about that revelation, in an attempt to influence their behavior. Comparisons across

treatment and control with this data measures the effect of the threat of monitoring via HFM on the MAOs.

The second round of data will consist of the second round of HFM data and the government MIS records as of the end of the second round of HFM. By comparing the extent of improvement between the first and second round of data across treatment and control, we measure the extent to which receiving the reports motivates the treatment MAOs to improve outcomes, and to which the information in the reports assists them in doing so. However, this effect is combined with a possible catch-up from the control group over time, revealing whether the reports actually improve outcomes or just shift receipt of cheques forward.

In the third round of data collection, we plan to conduct an endline survey data. The pre-analysis plan for that analysis will be filed later.

The data sources that we will use for the first and second round of analysis are the following:

1. Phone call data from first and second round of HFM
2. IVR call data from the first and second round of HFM
3. Government database of farmers registered for the scheme: contains name, district, mandal, village, phone number, and land size
4. Government MIS database: this records whether each beneficiary received a cheque and whether that cheque was encashed, as well as the date of those events
5. 2011 Indian census data at the block level

The data sources that we currently plan to collect and use in the third round of analysis are:

1. IVR data from endline survey
2. Phone survey data from endline survey
3. Field survey data from endline survey

C. Primary Outcomes of Interest from HFM calls and Government Records

PI1: A binary variable for whether the household received their cheque [HFM, admin]
A simple measure of program success is whether the household received their cheque. The MAO has control over this through scheduling of convenient cheque distribution days, ensuring that the cheques made it into the right hands, and organizing systems to distribute cheques to individuals that were absent on the day of distribution in their village.

PI2: A binary variable for whether the household encashed their cheque [HFM, admin]
The goal of the program is for households to encash the cheques and purchase agricultural inputs. This measures whether the household is able to do this. Key hold-ups include whether the bank was open for cheque encashment.

PI3: Beneficiary Satisfaction with the Implementation of Rythu Bandhu Scheme [HFM]
Satisfaction summarizes the response of beneficiaries to many types of problems and implementation successes that the MAO could affect. This has five responses: very dissatisfied, dissatisfied, neither satisfied nor dissatisfied, satisfied and very satisfied.

PI4: Corruption and Leakage [HFM]
During the HFM survey, we measure whether the respondent was asked to give a bribe in order to get or encash their cheque. By informing the MAOs and other government officials that we would be collecting beneficiary-level data on this form of corruption, this may increase their fear of exposure and reduce the extent of corruption

PI5: Speed of Distribution [HFM, admin]
As part of this program, all cheques were supposed to be distributed by May 20, 2018. We will ask respondents either whether they received their cheque by that date on the IVR calls and for the date of receipt on the phone HFM surveys. We will test whether treatment mandals had distributed a larger fraction of cheques successfully by that date as a binary measure of speed of distribution. This is important, since farmers needed the cash to purchase inputs for the upcoming growing season. We will also measure the date of cheque disbursement using the phone HFM and admin data, and compare the number of days from program start (May 8th) to date of cheque disbursement between treatment and control.[3]

D.Secondary Outcomes of Interest

This section lists outcomes that will not be the main focus of our analysis, but increase understanding of the results and support the primary analysis.

E.1. Reliability of Government Administrative Data
From the government of Telangana's MIS system, we get separate measures of whether each farmer (1) received and (2) encashed their cheque. This is based on reports from those distributing the cheques, and not subject to any form of biometric or similar authentication. Instead, agricultural extension officers were in charge of checking the identification of the farmers picking up the cheques and noting down who had picked up their cheques. The government did attempt to set up a form of validation (officials were supposed to take pictures of the farmers when they received their cheques), but the system worked imperfectly and is difficult to verify without a database of previous photos of farmers and substantial manpower to check them.

[3] In cases where individuals had not yet received their cheques, we will code the date of cheque disbursement as being the date of the survey for robustness.

As a result, it is possible that officials might misreport whether individuals received their cheques. For example, an official might report that a farmer received their cheque in order to make it appear that distribution in the area was more successful than it actually was. It is also possible that officials might collude with bank officials to encash cheques themselves, without giving them to beneficiaries.

We will test whether the government administrative data is reliable by comparing it to the HFM data. We will match the individual level responses from the HFM surveys to administrative data for that individual to check whether they are in agreement on whether the cheque was disbursed and encashed. We will similarly check whether the match rate differs between treatment and control to see whether mismatch is related to our treatment.

E.2. Additional Outcomes

During the HFM phone calls, two other outcomes will be collected that we list as secondary, rather than primary outcomes:

1. Did the farmer get the cheque at the *gram sabha* meeting in their village? If the individual did not receive it at this meeting, it suggests that they would have had to do more work to get their cheque, such as going to the tehsil office. It is a positive for implementation if a larger number of individuals received their cheques at the meeting.
2. Do they know of other people in their village who were supposed to receive cheques but did not? This question is intended to account for the fact that we only collect data from individuals who have phones. If implementation is substantially different for the population of individuals who lack phones, the HFM data will not capture that from direct reports. We thus also ask for indirect reports about others in order to capture problems in cheque distribution for the broader population.

E.3. Descriptions of Implementation

There are some features of the implementation that are unrelated to our treatment. For example, land records were updated prior to the announcement of the HFM system, and so the treatment will not have an effect on that. Nonetheless, it is useful to document the extent of such problems. We will provide summary statistics on:

1. Whether their cheque was the correct size; since this is based on land records, it should not be related to the treatment.
2. If they did not receive their cheque, what was the reason (if they know)?
3. If they received but did not encash their cheques, what was the reason?
4. What did they use the money for? The government intended that the money be used for purchase of farm inputs, but was unconditional, so could be used for anything.

E. Statistical Methods

F.1. Balance

At the mandal level, we will test for baseline balance on:

(i) 2011 census variables: % of households within 5km of a bank; % of irrigated land; % literate; % SC/ST. This will depend on whether or not we are able to match the 2011 census to current mandals, since the Government of Telangana recently changed the boundaries for a number of mandals.
(ii) Government administrative data: mean land size, median land size, 25th percentile of land size and 75th percentile of land size holdings within the mandal

F.2. Estimation

We will report ITT estimates, which compare average outcomes in treatment and control areas. Our primary outcomes are defined at the household level, which is the unit at which we will analyze them. Regressions will include fixed effects at the level of the randomization stratum and will be estimated using inverse sampling probabilities as weights. Standard errors will be clustered at the unit of randomization (MAO).

When analyzing the first and second rounds of data, we will run the specification:

$$Y_{ivmsd} = alpha + beta * treatment_{ivmsd} + delta_{sd} + gamma * X_{ivmsd} + epsilson_{ivmsd}$$

Where i is the individual, v is the village, m is the mandal, s is the stratum and d is the district. $delta_{sd}$ is a stratum fixed effect and X_{ivmsd} is a vector of baseline characteristics of the household that we observe in the administrative data (land size). We will also conduct randomization inference as a robustness check.

When analyzing the second round of data, we will include data from both rounds in the following specification:

$$Y_{ivmsd} = alpha + beta_1 * treatment_{ivmsd} + beta_2 * round2_{ivmsd} * treatment_{ivmsd} + beta_3 * round2_{ivmsd} + delta_{sd} + gamma * X_{ivmsd} + epsilson_{ivmsd}$$

Where round2 indicates that the observation is from the second round of data collection. We will test $H_0: beta_2 = beta_1$ in order to test if the treatment effect is the same across the rounds.

E.3. Robustness Checks

E.3.A. Comprehension of Treatment

One concern, particularly if the results are null, is whether this is due to misunderstanding on the part of treatment and control MAOs. For example, treatment MAOs may not full comprehend how the monitoring system works, or control MAOs may place a probability greater than zero of them being part of the treatment group. As a check of this concern, we are planning to run a survey of a sample of treatment and control MAOs to gauge: 1) whether they correctly perceived their treatment status; and 2) whether they understood the treatment. Results will be interpreted

in light of this analysis for robustness, such as rescaling estimates to account for the percent of mistaken treatment/control MAOs.

E.3.B. Tests for spillovers

One concern is that the estimates we generate are not generalizable to the case of full implementation of the HFM system due to spillovers across mandals. The main concern is that an HFM system reallocates the attention of those who supervise MAOs towards treatment mandals and away from control mandals. In this case, our estimates would be biased since the treatment is having a spillover effect on control mandals (reducing supervision) in addition to increasing upper level supervision of these mandals.

We will directly test for this effort by taking advantage of randomization of treatment intensity within existing bureaucratic units. Districts are divided into divisions, each containing between 3-10 mandals, which are overseen by an ADA. Within each of these divisions, there will be random variation in the number of treatment mandals. We will test whether in divisions that have more treatment mandals, there is decreased implementation quality in the control mandals (as would be consistent with a redirection of the attention of the ADA). To test this, we will run the following regression on the set of control mandals:

*Y_md = a + b * (number of treatment mandals in division)_md + (FE for total number of mandals in the division)_d + (district FE) + e_md*

Where the standard errors are clustered at the division level.

We will test whether the coefficient b is significantly different from zero. If so, and it is negative, this suggests that the attention of ADAs was diverted, meaning that our estimates may not generalize to the case of full implementation.

E.4. Heterogeneity analysis: (wealth index, whether own a cell phone)

We want to test whether effects of the treatment are heterogeneous depending on the individual characteristics of the beneficiary. We will include relevant interaction terms for these characteristics.

- Land size, as a measure of wealth. Since this distribution is right-skewed, we will take the natural log of land size. We will then recode the bottom 1% and top 1% of values to be equal to 1st and 99th percentile of land-holdings respectively.
- Mobile phone ownership - it may be that this improves outcomes for farmers with a phone, but worsens outcomes for other farmers, since bureaucrats prey on them. We cannot use this in the HFM data, since all surveyed farmers have phones, but can do so in the government administrative data (and will do so at endline).
- We would also like to test for heterogeneity with respect to measures of governance quality. Given the short-time frame, we are still investigating the availability of data on governance quality that could be used.

F. Randomization and Sampling

Our study population consists of all households eligible to receive the Rythu Bandhu benefits, i.e. all landowning farmers in the state of Telangana. Telangana contains 31 districts, but we exclude 1 from the analysis (Hyderabad) since it is urban and does not contain a significant number of program beneficiaries. We focus on a sub-unit of the district, the mandal, of which there are 584 in the state.

B.1. Mandal-level randomization

Randomization: In each district, we assign approximately 1 in 4 MAOs to treatment for a total of 120 treatment mandal agricultural officers. Our randomization is at the MAO-level because there are many cases where the mandal agricultural officer is responsible for overseeing multiple mandals. For policy relevant estimates, we do not allow within-MAO differences in treatment status across mandals -- in that case, the MAO may divert attention from treatment to control mandals, biasing the estimates.

We allocate the number of treatment mandals to a district based on the total number of MAOs in the district, in order to have a roughly equal proportion of MAOs across all of the districts. To calculate the number of MAOs to be sampled as treatment per district, we:

- For each district,take floor((number of MAOs within given district) / (total MAOs) * 120). This is the base number of MAOs to be sampled for that district. If a district is only assigned 1 MAO based on this formula, we round up to assign 2 base MAOs to treatment for that district (two cases). This assigns a total of 104 of the total of 120 treatment slots.
- We allocate the remaining 16 slots using probability proportionate to size sampling, where the population for the district under PPS is the remainder from [(number of MAOs within given district) / (total MAOs) * 120]. This means that those districts that were close to being assigned an additional treatment slot (those with reminder of close to one) have an appropriately high probability of being assigned a slot, while those who were not close to being assigned an additional treatment slot (those with remainder close to zero), have a low probability.
- As a result, all districts have the same number of expected treatment slots assigned to them as a function of their number of MAOs out of the total population of MAOs.

We now select mandals within each district. Due to a lack of mandal-level data at the time of randomization, it was not possible stratify the randomization using mandal-level characteristics. We stratify along the only mandal officer-level characteristics that were observed at the time of randomization: whether the MAO is in charge of multiple districts.

Within each district, we aim to create two strata (s_{MM} for MAOs who manage multiple mandals, s_{1M} for MAOs who manage a single mandal) with size approximately in proportion to that group's size out of total MAOs. We then select MAOs from each stratum in proportion to its fraction of the total number of MAOs. The process was:

1. Generate a unique list of MAOs within the district, as well as a count of how many mandals they manage
    a. Let N be the total number of MAOs in a district and N_{MM} be the number of MAOs who manage >1 mandal in the district. Let s be the total number of treatment slots assigned to the district. Define q^N as the quotient of N/s and r^N as the remainder of N/s.
2. If N/s < 2 (one case):
    a. Count entire district as one strata
    b. Randomly assign s MAOs to treatment, rest to control as a simple random sample.
3. If N >= 8:
    a. Randomly select r^N MAOs from the full list of MAOs via a simple random sample. These MAOs are assigned to control.
    b. The remaining $s*q^N$ MAOs will be divided into two strata. The first strata approximately corresponds to the set of MAOs who oversee multiple mandals, and the second strata approximately corresponds to the set of MAOs who oversee one mandal. However, the strata are sized to be evenly divisible by q^N in the manner described below to ensure that all MAOs in a district have an even probability of being assigned to treatment.
        i. Create an ordered list of the remaining MAOs. The first entries in the list will be those MAOs with multiple mandals, sorted in a random order, followed by the MAOs with only one mandal, sorted in a random order.
        ii. There will be a split point in the list after either entry q^N, $2q^N$, or $3q^N$. Mandals after the split are in the first stratum (s_{MM}), while those after the split are in the second stratum (s_{1M}). The split point is equal to the number of multiple mandal MAOs rounded to the nearest value of q^N, $2q^N$, or $3q^N$.
            1. Within strata s_{MM}, the number of MAOs assigned to treatment is $t_{MM} = \max(1, \text{round}(N_{MM} / (N/s))$
            2. Within strata s_{1M}, the number of MAOs assigned to treatment is $t_{1M} = s - t_{MM}$
            3. The size of each strata is
                a. $s_{MM} = q^N * t_{MM}$
                b. $s_{1M} = q^N * t_{1M}$
    c. Randomly assign t_{MM} (t_{1M}) MAOs from strata s_{MM} (s_{1M}) to treatment and the rest to control[4].

[4] An example helps to illustrate the process. Imagine a district with the following 9 MAOs, where the number in brackets indicates the number of mandals that they oversee:

(2) (1) (1) (2) (1) (1) (1) (1) (3)

For the sake of example, suppose that the district were allocated 4 treatment slots. Since 9 does not divide evenly into 4, we randomly select 1 mandal to be allocated into control. For the sake of the example, let's say this is a (1) MAO. We then sort the remaining MAOs into two strata as described, ordering them in a random order within multiple mandal and non-multiple mandal MAOs.We split into two strata as denoted by the line

(2) (3) (2) (1) || (1) (1) (1) (1) Control: (1)

Since there are two strata, we randomly select two from each stratum. This leads to:

The simulation graph shows the probability of each mandal being assigned to treatment. The variation is minimal and concentrated around a fixed value for all mandals. The only exception is one district (Warangal) for which the expected probability of allocation to treatment is 0.28, slightly above the rest. This is because the district only contains 7 MAOs, and so would only have been allocated one base MAO. Since we wished to ensure that all districts had at least two MAOs sampled, this slightly increases the likelihood of treatment MAOs coming from that district relative to its population of MAOs.

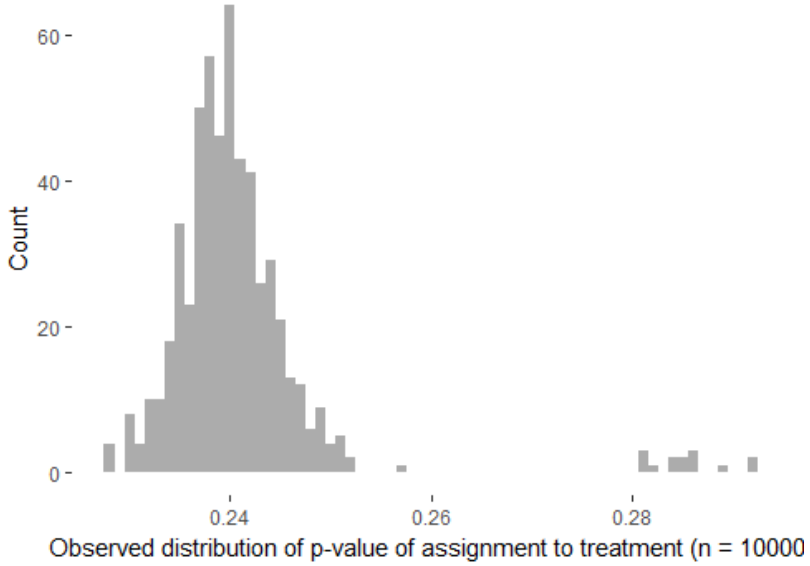

B.2. Individual Level Sampling for HFM calls

We draw a random sample of individual farmers whom the HFM system would call.

First Round of HFM calls in Treatment Mandals

In treatment mandals, the goal is to conduct 75 human-based phone surveys and 275 IVR-based phone surveys.

- We divide the farmers into five strata based on where they fall in the land size distribution of their village (0-20%, 21-40%, 41-60%, 61-80%, 81-100%). We then drop farmers who do not have phone numbers from the sample frame.[5]

Treatment: (2) (2) (1) (1) Control: (3) (1) (1) (1) (1)

Typically, one of the strata will contain both multiple mandal and single mandal MAOs, but this still decreases the odds of an unbalanced sample on this characteristic.

[5] There are some phone numbers that are linked to multiple farmers. In these cases. If the phone number is linked to an implausibly large number of farmers (more than 10), then it is dropped from the phone calling sample (2,805 numbers) since it is typically an indicator that the number belongs to a village level

- Using PPS sampling based on the population of farmers registered for RBS, we draw a random sample of 15 GPs. If fewer than 15 GPs, we select all of the GPs within the mandal
- . Within each stratum in sampled GPs, we randomly sample 2 households for a phone survey and 9 households for IVR surveys. Assuming a successful call completion rate of ½ for in-person calls and one-third for IVR calls, this will complete all of the targeted phone calls and 225 of the IVR calls[6]
- If there are fewer than 15 GPs in the mandal, within each strata we sample
    - 2*(15 - #GPs) farmers for phone calls
    - 9*(15 - #GPs) farmers for IVR calls
- For the remaining 50 IVR calls, we randomly select 30 farmers per stratum for calls.

<u>First Round of HFM calls in Control Mandals</u>

In control mandals, the goal is to conduct 25 human-based phone surveys and 75 IVR-based phone surveys

- We divide the farmers into 5 strata based on where they fall in the land size distribution of their village (0-20%, 21-40%, 41-60%, 61-80%, 81-100%, where there is an equal number of farmers in each stratum within each village). We then randomly sample 10 farmers from each stratum for phone calls and 45 per stratum for IVR calls. Since each stratum has an equal number of farmers, they all have an equal probability of being selected.

B.3. Reweighting

We save each household's ex ante probability of being included in the sample and will use these to weight estimators to make them representative of the overall frame of Rythu Bandhu beneficiaries. A farmer's sampling probability is the product of the probability that their district was selected, that their mandal was selected, that their village was selected and that the farmer was selected from its village.

official rather than an actual farmer. In the remaining cases, we call the number and match it to the record in the database using the name of the farmer on the call. For the purposes of sampling, the mobile numbers with multiple associated farmers are randomly assigned to the land size of one of those farmers, with equal probabilities associated to each entry.

[6] If the stratum has <11 farmers, then we sample all farmers from that stratum. We assign 2 farmers to phone surveys, as before. We then randomly assign the remaining farmers in that strata to IVR calls. If a stratum within a sampled village has no farmers after removing farmers with no phone numbers, then we do not sample farmers in that strata from that village. These cases are extremely rare, affecting less than 1% of the villages.