Targeting, Personalization, and Engagement in an Agricultural Advisory Service*

Susan Athey, Stanford GSB

Shawn Cole, Harvard University

Shanjukta Nath, University of Georgia

Jessica Zhu, PxD

October 15, 2024

Abstract

Information and Communication Technology is increasingly used to deliver customized information in developing countries. We examine whether individually targeting timing of automated voice calls increases engagement in an agricultural advisory service. We design, estimate, and evaluate a novel recommendation system that customizes contact times to individual characteristics. This generates significant gains, up to an 8% increase over the baseline pickup rate of 0.31. However, our on-policy estimated gains are lower than predicted by off-policy analysis. We show accounting for evolution of user preferences over time in off-policy estimation improves performance. We also demonstrate how this approach can be used to target vulnerable populations, and introduce a technique to quantify equity-efficiency trade-offs, and measure the social cost of resource constraints.

Keywords: Recommendation Systems, Agricultural Extension Services, Heterogeneous Treatment Effects, Causal Inference and Machine Learning.

JEL Classifications: O13, C90, Q16.

---

*athey@susanathey.com, scole@hbs.edu, shanjukta.nath@uga.edu, and jzhu@precisiondev.org. Authors are listed alphabetically.

The Golub Capital Social Impact Lab at Stanford GSB and the DFRD at Harvard Business School funded this research. The authors thank Sam van Herwaarden, Prasanna Kumar, Otini Mpinganjira, Tarun Pokiya, Aparna Priyadarshi, Theresa Solenski, and Hannah Timmis at PxD and Ayush Kanodia at Stanford University for their collaboration and support throughout the project. Harvard IRB reviewed and approved the data usage of this study (Protocol IRB24-0016). This research has been registered at AEA RCT registry (AEARCTR-0013161).

1 Introduction

Technology-enabled interventions can improve people's access to useful information, thereby increasing welfare and reducing inequality. The benefits of automated delivery of information have been shown in various sectors, such as public health, education, labor, and agriculture.[1] Information delivery through mobile phones may hold particular promise in developing-country settings, where access to the internet is still limited[2], and information barriers can be severe.

Automated delivery further enables the customization of both the content and the schedule for delivery. In this paper, we study a particular form of customization — the scheduling of delivery of information at a convenient time for the recipient — in the context of an agricultural advisory service in India that serves over one million farmers.[3] The service delivers weekly information to farmers via prerecorded telephone calls; the information is customized to the farmer's activities (e.g., geography, land type, seed variety, planting method) at a specific point in the growing season.

This project develops, implements, and evaluates a program that targets call times to farmers based on their observable characteristics and history of interaction with the service. Our primary outcome for farmer engagement is binary pickup when push calls are made by the service. We find substantial scope for gains from customization, with an 8% (Std. Err. = 0.94%) improvement in user engagement from a baseline of 31 percentage points (pp) relative to a non-optimized policy. If we were to implement the policy on the entire user base of 1.3 million farmers, this would translate to additional engagement for approximately 34,000 farmers (over the baseline non-optimized policy) with the service each week.

Our approach also allows us to identify efficiency-equity trade-offs in implementing targeted policies, and we show that a small sacrifice in overall expected outcomes can yield a significant increase in gender equity. The task at hand is to find the "best call time" for every farmer. This is a constrained optimization problem as there are technological limits on the number of calls that can be sent each hour from the service platform. Our setup involves 91 such hour-day combinations that comprise the treatment arms. By modifying welfare weights associated with subgroups, we can optimally allocate high-engagement slots

[1] The benefits of automated delivery are demonstrated in public health (Araya et al., 2021; Berman and Fenaughty, 2005; Ekeland et al., 2010; Knight et al., 2021; Lee et al., 2021; Voss et al., 2019), education (Agrawal et al., 2022; Hoxby, 2014; Rodriguez-Segura, 2022), labor (Dammert et al., 2013, 2015), and agriculture (Cole and Fernando, 2021; Fabregas et al., 2019).

[2] While the share of mobile cellular subscriptions reached above 90%, the percentage of population that use internet was still below 40% in lower- and middle-income countries in 2017 (Spielman et al., 2021).

[3] The advisory service served 1.3 million farmers at the time of this study, while its user base has grown over time. As of 2024, approximately 6.8 million farmers are registered with this service.

to targeted populations. For example, the gap between male and female engagement in the baseline is 3pp. We estimate targeted policies that can help reduce the gender engagement gap by 2.4pp (Std. Err. = 0.98pp, 80% reduction over the baseline difference) and quantify the "cost" of doing so in terms of the reduction in the total number of farmers reached.

We similarly show the value of targeting based on a proxy of need, focusing on farmers who lack smartphones (Giulivi et al., 2023; Ma et al., 2023). We develop and evaluate targeted policies that can preferentially allocate the farmers without access to smartphones to the high engagement call times using a welfare function that values the non-smartphone users 10 times relative to the smartphone users. This exercise helps clarify the potential benefits of investing in additional technological capacity, by identifying how many more non-smartphone and smartphone users could be reached.

These increased engagements with digital advisory service are the first crucial step in increasing farmers' knowledge and adoption of agriculture practices, which are the necessary conditions for enhancing agricultural productivity and livelihood. The importance of improving and retaining user engagement is critical for the success of digital interventions. This aspect has been studied in several domains such as marketing (Bruce et al., 2017; Hitsch et al., 2024; Yang et al., 2023; Yoganarasimhan et al., 2023), public health (Liu et al., 2023; Sadish et al., 2021; Yardley et al., 2016) and financial education (Blanco et al., 2023).

Digital advisory services have the ability to overcome the shortcomings of traditional in-person extension services and reach previously unconnected rural households with timely and customized information at low cost. While the efficacy of any particular program will depend on numerous factors (Abate et al., 2023; Spielman et al., 2021), such as the specific information gaps faced by farmers, the quality and relevance of information, and the information delivery timing and methods, a growing body of evidence finds such mobile services can improve agricultural practices, productivity, and profits (Arouna et al., 2020; Casaburi et al., 2019; Fabregas et al., 2019; Larochelle et al., 2019). In the Indian context, the value of information has been well documented in many areas, such as providing smallholder farmers with information on modern inputs and practices, weather forecasts, and output prices (Baul et al., 2024; Burlig et al., 2024; Cole and Fernando, 2021; Fernando, 2021; Giné et al., 2008; Jensen, 2007; Mitra et al., 2018). In all cases, reaching the farmer is an important and necessary first step.

Prior to our engagement with the NGO, calls were spread roughly across the week to manage limits on the number of calls that can be made by the platform in an hour. During that time, no efforts were made to match call times to user preferences or availability. In collaboration with the NGO, we changed the "default" policy from a purely ad hoc one to a purely random "uniform policy" in which users were assigned randomly with equal probability to each of the 91 potential calling times (treatment arms) within the week. This modification in the default policy is a straightforward change that requires minimal cost, effort, and no disruption for the platform. However, the uniform randomization of call times rules out any selection in assignment to different treatment arms and ensures equal representation of all farmer types across the 91 treatment arms.

This engagement data collected using uniform randomization helps to estimate the average treatment effect of different call times without any selection bias. The average engagement in the morning and evening hours is 1.3 pp (Std. Err. = 0.069 pp) and 1.8 pp (Std. Err. = 0.069 pp), respectively, higher than the afternoon hours. Moreover, we find high engagement hours during the weekends, which were far less used historically by the NGO to call farmers.[4]

We used data from the outcome of this "uniform randomization" to train a machine learning model, with the goal of maximizing user engagement while obeying bandwidth constraints, resulting in an "estimated optimal policy." Next, we conducted a prospective randomized evaluation, which compares the estimated optimal policy against the uniform policy by randomly assigning farmers to either their predicted optimal time (the targeted policy) or a time at random (the uniform policy). Using the data collected from these two policies allowed us to conduct two distinct types of evaluations. First, we conducted "on-policy" evaluations: since farmers were randomly assigned (as part of real-world service operations) to two groups, where one group received the targeted policy and the other the uniform policy, we were able to estimate the treatment effect of the targeted policy by comparing the sample mean of outcomes from the two groups of users. Second, we conducted "off-policy" evaluations, where we estimated the counterfactual benefits of alternative policies (policies mapping farmer characteristics to call times) that were not used in practice. Our off-policy evaluations took advantage of the fact that millions of calls were randomly assigned to call times under the uniform policy. However, for both the "off-policy" and "on-policy" evaluations, separate data sets were used for policy estimation and evaluation so that the benefits of the targeted policy are not overstated. Particularly, the off-policy evaluations are different from in-sample prediction, as we used cross-fitting in our off-policy evaluation. Cross-fitting splits the data into multiple folds and uses separate data sets for the estimation of targeted policy and evaluation of the benefits of the targeted policy (see

[4] Appendix Table A1 shows the distribution of calls across 91 hour-day call times historically under default ad hoc policy. The graphs show that a smaller proportion of calls were made during the weekend even though the service operated all 7 days of the week.

Section 6.1 for details).

We repeated the process of estimating optimal policies based on previous data and then implementing them in a randomized evaluation several times, each time updating the estimated optimal policy using the most recent data. We also provide new evidence on the reliability of off-policy estimates by comparing the gains estimated from off-policy estimates obtained at the point of policy design (using data until time t but with cross-fitting) to actual gains that resulted from on-policy evaluation in future weeks after time t.[5] This analysis highlights the extent to which changes in user behavior (driven, for example, by changes in preferences, circumstance, or other shocks) created differences between on- and off-policy estimates. Our off-policy estimates were 2.6pp (Std. Err. = 0.28pp) over the uniform policy group, but the on-policy estimates on future weeks were only 0.4pp (Std. Err. = 0.10pp).

The developing-country setting where we implemented our interventions introduced significant challenges, requiring us to use novel approaches to develop and evaluate recommendation systems.

First, the mission of the NGO with which we partnered is to improve the lives of the very poor. All of their communication occurs via telephone, both because many farmers have limited literacy and because few have access to the Internet. Although the user base of 1-1.3 million farmers is large relative to many economic field experiments, it is small relative to settings where recommendation systems, such as large online shopping platforms or popular apps, have been employed in practice. Relative to settings where policy estimation and evaluation have typically been rigorously analyzed, which usually involve a handful of treatment arms, our setting with many treatment arms (possible call times) introduces additional challenges. We demonstrate the benefit of deploying policies rather than using historical data for off-policy evaluation; even with a million users, on-policy evaluations have substantially more statistical power than off-policy evaluations due to the large number of treatment arms.

Second, because the agricultural advisory service, which is run in partnership between a nonprofit organization and the government, has significant bandwidth constraints, unable to call more than 70,000 farmers in any given hour, the development of targeted treatment assignment policies requires complex optimization, as an unconstrained algorithm would seek to schedule more than 70,000 farmers on high engagement hours.[6] This further intro-

[5]Cross-fitting refers to splitting the sample into K folds where a separate estimation is done for farmers in each fold using the data from remaining folds excluding the fold to which farmer i belongs. Moreover, estimated model parameters for every fold are saved and used as an input into policy evaluation.

[6]Note that there were 1.3 million total farmers in the sample at the time of the survey for whom calls

duces the possibility of trade-offs between equity and efficiency when determining whether to allocate high-engagement, scarce time slots to vulnerable groups. Our approach allows us to quantify the magnitude of these trade-offs for vulnerable groups such as female farmers or non-smartphone users.

Third, agriculture is a seasonal undertaking, and farmer behavior (or technology) may be subject to time-varying shocks; this can degrade the performance of recommendation systems in future weeks. We suggest two approaches to mitigate this cost using off-policy evaluations. First, we show that placing greater weight on more recent data than equally weighting training samples from a longer time series yields substantial gains. Modifying the weights of the training sample by proximity between the training and test sample shows gains in engagement between 2.0pp (Std. Err. = 0.59pp) and 2.7pp (Std. Err. = 0.49pp) on future weeks of data when we implement a targeted policy instead of uniform randomization. Second, the call center did two follow-up calls if the farmer did not pick up the call on the previous attempt. When we implemented the estimated optimal policies in practice, we only customized the time of the first attempt. In the presence of shocks, the first call time might fail, but we show using off-policy evaluation that if the advisory delivery system can potentially customize the time of the follow-up calls, further gains are possible. Additionally, we document using the follow-up calls on the "on-policy" that improvement in overall pickup is 0.8pp (Std. Err. = 0.10pp), which is twice the effect we see on the first call pickup. This supports the finding that follow-up calls can be used to improve the performance of targeted policies in the presence of shocks.

Although this paper evaluates the value of personalized policies in terms of an important, "advisory content delivered," given the short time frame of the experiment, we were not able to assess the gains on downstream outcomes such as agricultural yields and farmer adoption of recommended practices. Nevertheless, a rigorous impact evaluation conducted on this digital advisory service shows that providing farmers with the advisory service, on average, increases their agricultural knowledge and adoption by 0.1 standard deviations, raises productivity by 1.74% and total production by 4.12%, and reduces their likelihood of experiencing severe rice crop loss by 10% (Cole et al., 2024). Impacts are larger among farmers who are more engaged with the service. Moreover, a canonical meta-analysis computes that digital interventions in agriculture and increased access to information have improved downstream outcomes such as agricultural yields by 4% (Fabregas et al., 2019).

had to be scheduled for at least 2 to 3 messages each week of the experiment. Additionally, each farmer gets 2 follow-up calls if they fail to pickup the call in the first attempt.

2 Related Literature

Several previous studies have used historical data, either experimental or observational, to estimate personalized (targeted) treatment assignment policies and evaluate the counterfactual benefits of implementing them. An early paper in this literature is Ascarza (2018), which applies machine learning methods to data from a randomized experiment to compare the performance of two alternative approaches to targeting. A traditional approach to targeting (often used in the absence of data from randomized experiments) prioritizes individuals for treatment based on the predicted baseline value of the outcome of interest, in this case, customer churn. Ascarza (2018) quantifies the benefits of targeting instead based on individual-level estimates of treatment effects derived from experimental data. Other papers that compare targeting based on estimated treatment effects to alternative targeting rules such as targeting based on predicted baseline outcomes include Athey et al. (2023b) in the context of text-message nudges to students filling out financial aid forms, Devriendt et al. (2021) in marketing, Inoue et al. (2023) in health, Olaya et al. (2020) in education, and Haushofer et al. (2022) for a cash transfer program in a developing country setting.

In addition to the above papers, several recent papers have focused on prioritization by treatment effects and evaluated the benefits of personalization counterfactually. For instance, Yoganarasimhan et al. (2023) uses data from a randomized experiment to estimate personalized policies using off-policy evaluation methods in an application designed to determine the optimal trial period to offer customers a software service. Similarly, Hitsch et al. (2024) uses data from randomized experiments to compare different targeting policies for catalog mailing. Yang et al. (2023) moves beyond targeting short-run outcomes and develops new methods to design and implement targeted policies for long-run outcomes. Our paper follows a similar approach but also builds on this literature in that we conduct a series of experiments. Prior to each experiment, we use historical data (observational at the start, and in later rounds experimental) to conduct estimation and counterfactual evaluation of personalized treatment assignment rules.

However, in addition to off-policy evaluations, we also deploy our estimated targeted policy in subsequent weeks. Very few studies are able to evaluate the performance of targeted policies by deploying them in practice, enabling the use of "on-policy" evaluation (using data from the real-world implementation of the targeted policy) in addition to the counterfactual "off-policy" evaluation. Three papers that estimate, implement, and evaluate targeted policies are Yang et al. (2023), Dubé and Misra (2023) and Simester et al. (2020b). Each of these papers runs a sequence of two experiments, where data from the first experiment is used to estimate a targeted treatment assignment policy and generate counterfactual estimates of its benefit. In the next step, the authors run an experiment where they deploy the estimated policy in practice, evaluating the benefits on-policy. Yang et al. (2023) use this approach to target discounts to digital subscribers of Boston Globe. Dubé and Misra (2023) study the implications of personalized pricing on demand and consumer welfare. Lastly, Simester et al. (2020b) evaluates the benefits of targeting on customer membership for a large US retailer.

Similar to these studies, we run multiple experiments, in our case testing four different personalized policies. We identify discrepancies between estimated (off-policy) and subsequent actual (on-policy) performance, showing that they arise at least in part due to a changing environment and proposing approaches to mitigate the problem. Two of the above-referenced studies, Yang et al. (2023) and Simester et al. (2020b), also address the issues with a changing environment. Simester et al. (2020b) evaluates the performance of targeted policies using data from two experiments conducted 6 months apart. The first experiment (training dataset) provides the training data to estimate 7 targeted policies using different machine learning methods. These targeted policies are deployed in the second experiment (validation dataset). The second experiment is used to evaluate how differences (seasonality, covariate shifts, loss of information due to data aggregation) between the training and validation data can result in the poor performance of targeted policies in the validation dataset. Yang et al. (2023) considers the possibility of a changing environment in their method for policy estimation, and their approach includes ongoing randomization and adaptive policy learning that enables the targeted policy to be updated in response to environmental changes.

In addition, the context of our study presents us with new challenges regarding capacity constraints on the number of calls that can be made from the call center in an hour. Historically, practitioners and academics have considered the problem of prioritizing a costly treatment to a subset of users (e.g., Ascarza (2018); Neslin et al. (2006)). In contrast, we solve for the optimal allocation of scarce slots across the entire population. Resources are often constrained in developing-country settings, a key difference from such technologies in developed countries. Consequently, our optimization approach for personalized policies requires addressing this added challenge due to our setup. Ours is one of very few research papers that augments recommendation systems and personalized policies (typically developed in settings without constraints, e.g., Athey et al. (2024); Bodapati (2008); Rafieian and Yoganarasimhan (2021); Zhou and Zou (2023)) by incorporating constrained optimization in a real-world application, with over a million customers and 91 treatments and directly measuring its impact. We use a two-step approach, using machine learning methods on the randomized data with the outcome model to predict pickup for each farmer in each of the 91 call times. In the second step, we use these predictions along with call time constraints (91 constraints, one for each call time as the number of calls cannot exceed the pre-existing capacity limit of the call center) and the farmer constraints (one for each farmer as each farmer should be placed in only one of the call times) to set up a mixed integer programming problem and solving for it. In this regard, we add to very limited work on incorporating constraints for targeted policies in a real-world setting. An exception is a recent work by Lu et al. (2023) where the authors use primal-dual hybrid gradient linear programming methods to incorporate fairness or volume constraints in estimating targeted policies.

Constrained optimization for personalization allows us to study the equity-efficiency trade-off and develop policies to improve engagement outcomes for vulnerable groups. This relates to the fairness discussion in machine learning (Athey et al., 2022; Beretta et al., 2019; Rambachan et al., 2020). We develop methods that can help the agricultural advisory quantify the social returns to investing in greater bandwidth, as well as quantify the differential impact on female and poorer subgroups. The constrained optimization methods for personalization used in this paper can be applied to several problems in marketing, such as fairness concerns in promotion campaigns, advertisements (Friedman et al., 2023; Li et al., 2024), risk assessment by insurance companies and banks, as well as targeting with capacity constraints due to limitations on technology or supply.

Several of the above-cited studies provide evidence of the value of personalized policies in marketing applications (Ascarza, 2018; Hitsch et al., 2024; Simester et al., 2020a; Yoganarasimhan et al., 2023). Our study has broader implications for marketing activities that rely on attracting user attention. Moreover, we focus on a social impact application in a developing country. The limited evidence on the value of targeted treatment assignment rules (Agrawal et al., 2022; Athey et al., 2023a) in developing country setup studies delivery techniques that require Internet access. We contribute findings for a technology (automated voice calls) that can serve the poorest billion individuals currently lacking internet access in a developing country. Finally, our focus on constraints and fairness can have significant implications for managers working on social impact applications.

3 Context

Improving agricultural practices is a key strategy to reduce poverty, promote food security, and address environmental concerns (Foster and Rosenzweig, 2010; Takahashi et al., 2020). A primary roadblock in adopting innovative agricultural practices is the limited access to information by smallholder farmers(Dzanku et al., 2020; Magruder, 2018). Eliminating this information gap is of paramount interest to policymakers and practitioners. In the Indian context, information constraints have been shown to be binding for farmers across a variety of topics such as adopting modern inputs and practices, mitigating climate shocks, negotiating for harvest sales (Baul et al., 2024; Burlig et al., 2024; Cole and Fernando, 2021; Fernando, 2021; Giné et al., 2008; Jensen, 2007; Mitra et al., 2018).

While many developing countries have invested heavily in agricultural extension services, their reach is often limited, and the empirical evidence on their efficacy is quite mixed (Abate et al., 2023; Anderson and Feder, 2004; Spielman et al., 2021). India faces such challenges. The rural population comprises 64% of the total population in India, and the majority (about 90%) of the poor reside in rural areas.[7] The Government of India operates a pluralistic extension system with 90,000 extension agents (Swanson and Davis, 2014). However, less than 6% of farmers report having received extension services in the past year in a survey with cotton farmers (Cole and Fernando, 2021). In the state where this study took place, 4,900 village agricultural workers, agricultural supervisors, and Assistant Agricultural Officers (AAO) serve eight million farmers.

A critical development in the past decade affecting the agricultural sector in developing countries has been the widespread availability of low-cost telephone services. Technology has opened up new opportunities for sharing information with farmers (Aker, 2011; Aker et al., 2016; Fabregas et al., 2023). The development of machine-learning techniques offers an even greater opportunity, potentially enabling agricultural advice to be customized, in an automated fashion, for millions of farmers. This paper adds to the nascent literature on bringing such techniques to large populations in a developing country (Agrawal et al., 2022; Athey et al., 2023a).

To investigate the impact of service customization and targeting, we conducted a multistage experiment with a phone-based agricultural extension service in India. This digital extension service, launched in 2018, has been developed and implemented by NGOs in collaboration with an Indian state government.[8] By the end of 2021, it served 1.3 million

[7] The percent of the population in rural areas in India is for 2022 from the World Bank's DataBank, and the fraction of poor corresponds to the 2022 Multidimensional Poverty Index (MPI) report by the United Nations Development Program (UNDP).

[8] Specifically, this digital extension service has been developed and implemented by Precision Develop-

smallholder farmers throughout the state with a two-way, mobile-phone-based platform and a live call center. This service has expanded rapidly, and as of 2024, it has 6.8 million enrolled farmers.

The service provides customized advisories on 21 crops, livestock, and fisheries using farmer covariate information (i.e., language, location, crop, water management) and agricultural data (i.e., weather forecast, market information, pest/disease outbreaks). Users of this service receive agricultural information through three channels: 1) weekly interactive voice response (IVR) calls that provide farmers customized farming advisories timed to the crop calendar (Outbound Calls); 2) an IVR platform that farmers can call in to listen to content from an advisory library and record their questions; and 3) a call center where farmers can call in and ask agricultural-related questions. Questions are answered by local agronomists, who send recorded answers within 48 hours.

The experiment was conducted over six weeks in October and November 2021. The weeks of this multistage experiment are defined in Table 1. Every week, we sent an agricultural advisory push call to nearly 1 million farmers, prioritizing advisories on rice, one of the most important staple crops in this Indian state. Appendix A shows example scripts of agricultural advisory messages sent to farmers.[9]

4 Setup

In this section, we first provide a high-level overview of the different data collection methods used during the six weeks of the experiments. Next, we introduce notation that will be needed to describe our experiments and analysis more formally.

As described in the introduction, two data collection methods were used for this experiment. The data collection method determines the probability (μ_{ij}) that farmer $i \in \{1, 2, .., N_t\}$ is called in day-hour block $j \in \mathcal{J} = 1, .., 91$. The first data collection method, which we refer to as "uniform randomization," assigns each farmer i to each of 91 call times, denoted j, with equal probability, so that $\mu_{ij} = 1/91$ for all i and j. The second data collection method, which we call a "targeted policy," is deterministic and depends on the observable characteristics of a farmer x_i. The targeted policy π is a mapping from farmer covariates x_i to treatment arms j. Thus, the probability that farmer i is allocated to call

ment and the Abdul Latif Jameel Poverty Action Lab in partnership with an Indian state government, with support from the Bill & Melinda Gates Foundation.

[9]We provide example advisory messages for two types of messages in Appendix A. One is an advisory message on pest management, and the other is an advisory message on basal fertilizer application for transplanting.

Table 1: Experiment Weeks in October and November 2021

Week No.	Date	Uniform Randomization		Targeted Policy	
		Sample Name	Sample Size	Sample Name	Sample Size
1	Oct 5-10*	$N_{1,\mathcal{U}}$	881,891	–	–
2	Oct 18-24	$N_{2,\mathcal{U}}$	616,656	$N_{2,\hat{\pi}_A}$	265,188
3	Oct 25-31	$N_{3,\mathcal{U}}$	879,109	–	–
4	Nov 1,Nov 4-7*	$N_{4,\mathcal{U}}$	707,644	$N_{4,\hat{\pi}_B}$	167,995
5	Nov 8-14	$N_{5,\mathcal{U}}$	624,801	$N_{5,\hat{\pi}_C}$	234,276
6	Nov 17-23	$N_{6,\mathcal{U}}$	587,608	$N_{6,\hat{\pi}_D}$	227,386

Notes: *In these weeks, data was collected for only the indicated days of the week and not the entire week. $N_{t,\mathcal{U}}$ is used to denote the sample that is allocated to uniform randomization in week t. $N_{t,\hat{\pi}_e}$ is the sample that receives targeted policy $\hat{\pi}_e$ where $e \in \{A, B, C, D\}$ and t denotes the week. Our sample includes 880,000 farmers for whom we observe a complete set of covariates. The sample varies from around 800,000-880,000, depending on which farmers are designated to receive messages each week.

time j is 1 if $\pi(x_i, \cdot) = j$. During the course of this study, we implemented several different targeted policies. In addition, in selected weeks, we implemented a higher-level experiment, where eligible farmers were randomized (with constant probability) into either the uniform randomization data collection or a targeted policy. Different targeted policies were used in different weeks. Data collected using the uniform randomization method from previous weeks was used to estimate a model and develop a targeted policy for the subsequent week. Figure 1 shows the data collection process for every week of the experiment.

Formally, the experiment was conducted over six weeks $t \in \mathcal{T} \equiv \{1, 2, 3, .., 6\}$. Each week of data is a sample N_t of farmers drawn from the population. N_t consists of those farmers who have a complete set of covariates and have signed up to receive messages assigned for week t. The treatment is call time, which is a combination of hour and day of the week. The call center operates from 8 AM to 9 PM, seven days a week, so there are 91 hour-day call times. The experiment is conducted on a weekly basis as new crop advisories are provided to the farmers every week through the extension service. Every farmer $i \in \{1, 2, .., N_t\}$ can receive only one of the individual treatments $j_{it} \in \mathcal{J} \equiv \{1, 2, .., 91\}$ in each week. The treatment J is a random variable with support in $\mathcal{J}$ and $J \sim F_J$. Note that j is a realization of J.

The outcome is binary pickup dummy Y_{it}^{obs}. Let $Y_{it}(j)$ be the potential outcome for farmer i in week t and call time j. The observed outcome can therefore be written in terms

Figure 1: Two Data Collection Methods

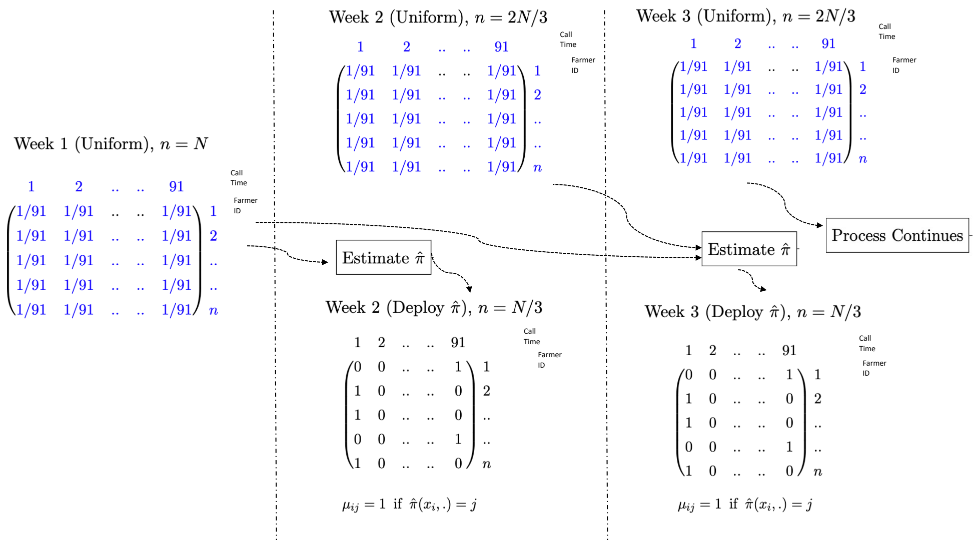

Notes: The figure illustrates the two different data collection methods. We start with uniform randomization between 91 treatment arms. This data is used to estimate the targeted policies. In subsequent weeks the targeted policies are deployed (our second method of collecting data).

of potential outcomes as $Y_{it}^{\text{obs}} = \sum_{j=1}^{|\mathcal{J}|} \mathbb{1}(J_{it} = j) Y_{it}(j)$. Also, Y_{it}^{obs} is a random variable with support in $\mathcal{Y} \equiv \{0, 1\}$. $x_i \in \mathcal{X}$ is the vector of observed covariates for individual i. This vector does not vary with time, so we do not have the t subscript. We let F_X denote the distribution of X with support in $\mathcal{X}$. Each week, the observed covariates are determined before an intervention is assigned. Table A1 shows the set of observed covariates. It includes gender, access to irrigation, land size, smartphone ownership, district of residence, and historical engagement with the service prior to the start of our experiment.

The potential outcome for farmer i in call time j for week t, $Y_{it}(j)$, follows a Bernoulli distribution with probability of pickup given by $\mu(x, j, t)$. Therefore, $\mu(x, j, t) = \mathbb{E}_Y[Y_{it}(j)|x]$.

In order to construct an estimator for μ, we need to specify a prediction model m (e.g., a LASSO regression model with a given specification), where we let $\mathcal{M}$ denote the set of possible models, and a training dataset $\mathcal{S} \in [\mathcal{Y}, \mathcal{X}]^n$. Then, we denote an estimator by $\hat{\mu} : \mathcal{X} \times \mathcal{J} \times \mathcal{T} \times \mathcal{M} \times [\mathcal{Y}, \mathcal{X}]^n \rightarrow [0, 1]$. Given a ML model $m \in \mathcal{M}$ and a training dataset $\mathcal{S}$, $\hat{\mu}(x, j, t; m, \mathcal{S})$ is the estimate of $\mu(x, j, t)$.

4.1 Targeted and Uniform Policies

Several policies $\pi : \mathcal{X} \rightarrow \mathcal{J}$ were estimated, deployed, and evaluated in this project. In most scenarios in this project, we value reaching each farmer equally and do not put welfare weights on the outcome. The population object of interest is therefore defined as:

$$V(\pi, \mathcal{S}^{\text{eval}}) = \mathbb{E}_{X\sim F_X}[\mathbb{E}_Y Y_t[(\pi(X))]]$$

where $\mathcal{S}^{\text{eval}}$ is the evaluation sample. Class Π is to denote the constraints used in this experiment. For this paper, it encodes budget constraints.[10] Therefore, the optimal policy π^* is $\pi^*(\Pi) = \arg\max\{V(\pi, \mathcal{S}^{\text{eval}}) : \pi \in \Pi\}$. This paper uses machine learning models to estimate targeted policies. The goal is to learn about the best call times for every farmer in the sample using predicted pickup rates. The estimator of π^* is denoted as $\hat{\pi}(X, t, \Pi; m, \mathcal{S})$. This estimator is a function of covariates x and is parameterized by the week t, class Π of policies, a ML model m, and the training data $\mathcal{S}$ used to estimate the parameters of the model. It returns the best call time based on $\hat{\mu}$. Hence $\hat{\pi} : \mathcal{X} \times \mathcal{T} \times \mathcal{P} \times \mathcal{M} \times [\mathcal{Y}, \mathcal{X}]^n \to \mathcal{J}$ where,

$$\hat{\pi}(x; t, \Pi; m, \mathcal{S}) = \arg\max_j \hat{\mu}(x; j, t; m, \mathcal{S}), \quad \text{s.t. } \pi \in \Pi$$

We call the estimated policies $\hat{\pi}$ "targeted policies" throughout this paper. The population object of interest for the targeted policy is the value of the targeted policy evaluated on an evaluation sample S^{eval} $(V(\hat{\pi}, \mathcal{S}^{\text{eval}}))$. Therefore $V : \Pi \times \{\mathcal{Y}, \mathcal{X}\}^n \to \mathcal{R}$ and

$$V(\hat{\pi}, \mathcal{S}^{\text{eval}}) = \mathbb{E}_{X\sim F_X}[\mathbb{E}_Y Y_t[(\hat{\pi}(X, \cdot))]]$$

As illustrated in Figure 1, dividing the population into a targeted policy and a group whose call time is uniformly randomized serves to evaluate the targeted policy against a "control group" (uniform randomization), but the control group itself generates data that can be used to evaluate counterfactual policies and help design subsequent targeted policies.

In order to motivate the definition of the uniform policy, we first define a "fixed time policy" π^j where every farmer is called during call time j. Formally, $\pi^j(x) = j \quad \forall x$. The uniform call time policy is explained using the following steps. For farmer i, step 1 is to draw a random variable J_{it} from a discrete uniform distribution $\mathcal{U}\{1, 91\}$. Step 2 is if $J_{it} = j$, then call farmer i in call time j. The population object of interest for the fixed call time policy is

$$\begin{aligned} V(\pi^j, \mathcal{S}^{\text{eval}}) &= \mathbb{E}_{X\sim F_X}\mathbb{E}_Y(Y_t(\pi^j(X))) \\ &= \mathbb{E}_{X\sim F_X}(\mu(X, j, t)). \end{aligned}$$

[10] In Section 7.4, we incorporate weights into this population object of interest to discuss the equity efficiency trade-off.

Consequently, the population object of interest for the uniform policy is

$$
\begin{aligned}
\bar{V}(\mathcal{U}, \mathcal{S}^{\mathcal{U}}) &= \mathbb{E}_{X \sim F_X}[\mathbb{E}_Y[\mathbb{E}_{J \sim \mathcal{U}}[\sum_j (Y_t(\pi^j)\mathbb{1}(J = j))]]] \\
&= \sum_j \mathbb{E}_{X \sim F_X}[\mathbb{E}_Y[Y_t(\pi^j)\mathbb{P}(J = j)]] \\
&= \sum_j \mathbb{E}_{X \sim F_X}[\mu(X, j, t))\mathbb{P}(J = j)] \\
&= \frac{1}{91}\sum_{j=1}^{91} V(\pi^j)
\end{aligned}
$$

where $S^{\mathcal{U}}$ is the evaluation sample for the uniform policy.[11] In this paper, the data collected using uniform randomization is used as a baseline as the uniform randomization data was collected at the same time as the outcome data for target policy. We do not use the historical data as a baseline because it was collected before the randomization and using it as baseline will not control for the temporal shocks to the outcome variable across different time periods. Moreover, Figure A1 suggests that very few calls were made during the high-engagement weekend hours. This limits the possibility of having enough farmers with varied characteristics across all of the 91 treatment arms.

5 The Effect of Call Times on Outcomes: Evidence from Uniform Randomization

As described in Figure 1, in this project, uniform randomization was used as a data collection method for a randomly selected set of farmers. Although the project's ultimate goals include estimating and evaluating alternative targeted policies, we motivate this work by reporting findings from the uniform randomization sample about the variation in engagement across 91 call times. These findings showcase the importance of call time as well as the heterogeneity of preferences across demographic groups.

Figure A2 shows farmers' engagement with the service. The outcome variable is a binary outcome representing whether a farmer picked up the first call.[12] The average pickup rate is around 31%, but there is substantial variation in pickup rates at different points of time

[11] In this paper, $\bar{V}$ is used to denote the value of a random policy like uniform policy, and V is used for the value of a deterministic policy like the targeted policies estimated in this paper.

[12] In most analyses in this paper, we focus on the pickup of the first attempt of the call, while each call is attempted up to three times. In Section 7.3, we discuss the follow-up calls.

during the week. This suggests that evening and morning hours have relatively higher pickup than afternoon hours. We demonstrate the magnitude and statistical significance of these differences in an abbreviated manner by showing the mean pickup rates for morning, afternoon, and evening times and comparing afternoon pickup rates with morning and evening pickup rates separately. Table 2 shows the average pickup in the afternoon is 1.3 (1.8) percentage points lower than the morning (evening).

The uniform randomization data helps us identify several hours during the weekends that are high engagement hours, especially the evening hours. Historically, the agricultural advisory service did not schedule many calls during the weekends even though it operated all seven days of the week. Hence, it had inadvertently failed to take advantage of very high engagement hours.

Table 2: Variation in Engagement by Call Times

	Morning	Afternoon	Difference
Mean	0.320	0.307	0.013
Std. Error	[0.00051]	[0.00046]	[0.00069]
	Evening	Afternoon	Difference
Mean	0.325	0.307	0.018
Std. Error	[0.00052]	[0.00046]	[0.00069]

Notes: We pool the data collected using uniform randomization in weeks 1, 2, and 3 of the experiment (see Table 1 for sample size). Morning hours are between 8 AM to 12 PM, afternoon hours are between 12 PM to 5 PM, and evening hours are between 5 PM to 9 PM.

Next, we present the socio-demographic characteristics of the sample and variation in engagement for different farmer subgroups. Table A1 provides summary statistics on farmer characteristics. About 18% of the sample are female farmers. Around 37.4% of the farmers are smartphone users, and a little over 44% use irrigation systems on their land. Additionally, the residential districts for the farmers are provided in the covariates data.

Previous research has shown substantial gender gaps in access to technology in agriculture (Owusu et al., 2018; Quisumbing and Pandolfelli, 2010). Consistent with the literature, we observe a persistent but time-varying gap in engagement rates for female farmers (0.29) relative to male farmers (0.32), as shown in Panel (a) in Figure 2. The graph also suggests that the high pickup times often overlap for male and female farmers. The overlap hours have an important implication for policy design, as simply maximizing overall pickup rates (equal weighting) could result in giving the most valuable time slots to men; alternatively, placing a higher weight on reaching women could preserve some of the higher-value slots for women. This could help reduce the gender gap in engagement with the service.

We further explore the heterogeneity in farmers' engagement by their access to information and wealth levels. Our first result is that the average pickup rate among smartphone users is 4.1 percentage points lower than among non-smartphone users. Figure A3 shows this comparison of engagement by smartphone ownership.

We then examine the variation in pickup rates by the distribution of land size, which serves as a proxy for wealth. We divide the sample of farmers into deciles based on their land size. Panel (b) in Figure A3 shows that farmers at the high end of the land size distribution have lower engagement than the farmers with smaller land sizes. These figures together suggest that non-smartphone users and poor farmers are more engaged with the service, consistent with the hypothesis that they have limited outside options to access farming-related information.

We additionally discuss some crucial features of the farmers based on their historical engagement behavior. We have data on farmers' historical engagement with the agricultural advisory service from the initiation of the service (July 2018) till right before our experiment (September 2021). We also use historical engagement data to estimate targeted policies. Figure A4 shows the distribution of pickup across the 91 call times in the historical data.

The farmers in the experiment sample registered with the agricultural advisory service at different points in time. Panel (a) in Figure A5 shows the distribution of the length of the total time with the service. The mean farmer duration with the service is 516 days. Panel (b) shows the variation in the month when they started receiving calls from the extension service. High enrollments are observed for the months of April, September, and December. Since the experiment was conducted in the months of October and November 2021, it is possible that farmers who signed up with the service at different points in time value the information sent out during the months of the experiment differently. Finally, we show the variation in engagement during the weeks of the experiment by their duration with the service. Figure A5 suggests the possibility of information fatigue directly related to the time the farmers have been associated with the service. We define a new farmer dummy, which takes a value of 1 if the farmer spent less than the mean duration with the service and 0 otherwise. The engagement of new farmers is consistently higher than the engagement of the old farmers (Panel (c) in Figure A5).

Finally, the assignment of calls to hours in the historical data followed a changing set of ad hoc rules, resulting in a non-uniform distribution of calls across the 91 treatment times. Figure A1 illustrates this separately for each historical year. In 2018, no calls were made on Tuesdays, and very few calls were made on Sundays. In 2020 and 2021, very few

Figure 2: (a) Pickup by time and gender, (b) Highest Pickup arms for female farmers and (c) male farmers in the dataset.

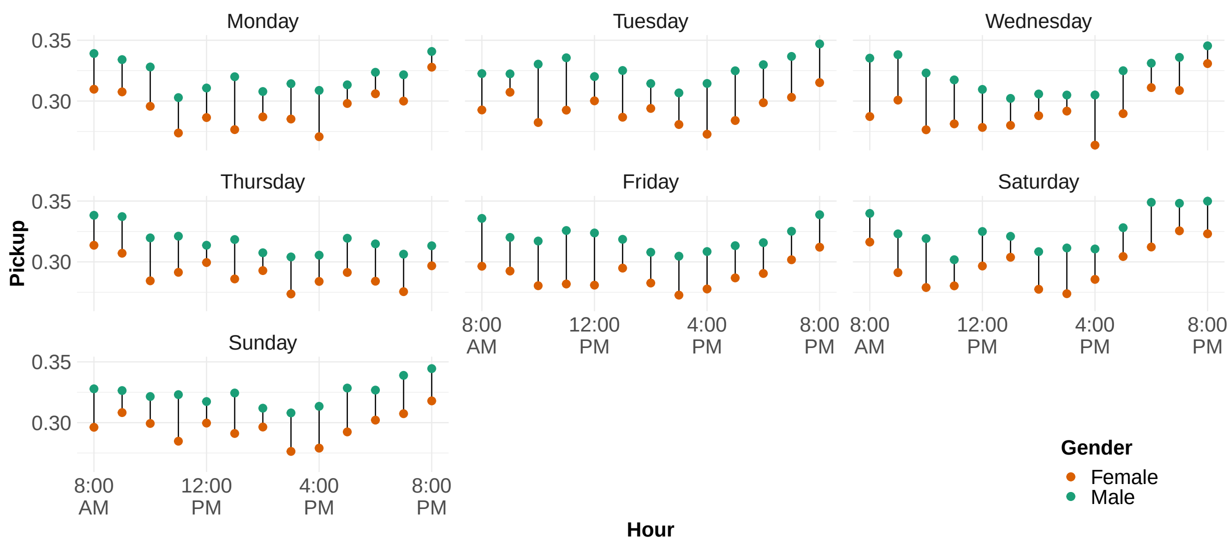

a. Variation in Pickup by Gender for 91 Treatment Arms

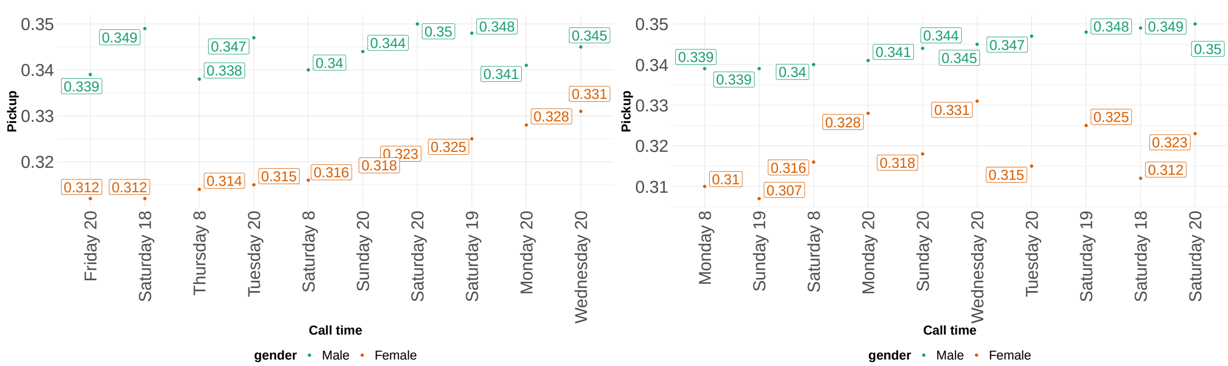

b. Top 10 Buckets for Females

c. Top 10 Buckets for Males

Notes: We pool the data collected using $\mathcal{U}$ in weeks 1, 2, and 3 of the experiment. Panel (a) illustrates the gender gap in engagement over the 91 call times. Panel (b) shows the most popular call times for female farmers arranged in increasing order of popularity, with the pickup for male farmers for comparison. Panel (c) shows the most popular call times for male farmers arranged in increasing order of popularity with the pickup for female farmers.

calls were made on Sundays, especially during the evening. This highlights the benefits of uniform randomization of call times to ensure that there is sufficient data to learn about the relationship between farmer characteristics and engagement across all call times.

6 Method for Estimating and Evaluating a Targeted Policy

As Figure 1 outlined, this project had multiple stages. In week 1, training data was collected using $\mathcal{U}$. In week 2, the uniform randomization data from week 1 was used to construct and deploy a policy in a randomized controlled trial, comparing the targeted policy to the uniform policy. In each subsequent week, data from the previous weeks' uniform policies were used to design a new targeted policy, which was again evaluated against the uniform policy.

6.1 Estimate $\hat{\pi}$ Using Uniform Randomized Data

This section describes how targeted policies were estimated. First, we selected an estimator for the outcome model. With sufficient data, it would be possible to separately estimate $\mu(\cdot, j, t)$ for each treatment arm j and week t. However, even though we have approximately 880,000 farmers in the sample with 91 treatment arms and many covariates, it was efficient to pool the data into a single model and to use data-driven model selection. Among many reasonable machine learning models (e.g., LASSO, random forest) or matrix factorization methods (e.g., recommendation systems), we chose LASSO regularized regression due to a large number of potential treatments and covariates. Our regressors include treatment indicators for each call time j, covariates x_i, and interactions of those covariates with treatment indicators.

Because the outcome model will be used to select treatments, it is important that our estimator $\hat{\mu}(X, J, t; m, \mathcal{S})$ yields useful estimates of the difference in expected outcomes across treatment arms. Since our training data $\mathcal{S}$ has uniform randomization, we do not need to be concerned that farmers assigned to different treatment arms are different in some unobserved way; the randomization ensures independence ($\{Y_i(j)\}_{j=1}^{J} \perp\!\!\!\perp J$) and thus unconfoundedness. Formally, the estimator takes the form:

$$\text{logit}(\mu_{ij}) = X_i\beta + \sum_{j=1}^{91} \delta_j J_i + \sum_{j=1}^{91} \gamma_j X_i J_i \tag{1}$$

where μ_{ij} is the probability of call pickup, X denotes the covariate matrix, and J denotes the treatment dummies. The objective of LASSO is to minimize the following:

$$-l(\theta, X, T) + \lambda\|\beta\| + \lambda\|\gamma\|, \tag{2}$$

where $\theta = (\beta, \delta, \gamma)$ and l denotes the log-likelihood of the outcome. The selection of the regularization parameter λ is done using cross-validation. We consider the full training data and split it into k folds. The optimal λ_0 minimizes the average mean squared error across the k folds. The mean squared error is computed as the squared difference between the observed and predicted outcome.

The model incorporates the main effects and the covariates' interaction with the treatment dummies. Lastly, we do not penalize the coefficients on the treatment dummies. We also estimate a few variations of the above specification. For instance, a model with a polynomial function for the treatment dummies is estimated in Appendix F.1. We also estimate an additional specification where the penalty on the regularization parameters varies based on the order of the interaction terms. We call this specification the hierarchical LASSO. The interaction effects are likely smaller in models with two-way and higher-order interactions than the main effects. Hence, the penalties are allowed to increase with the degree of the interactions. Appendix F.2 discusses this variation and result.

Below, we evaluate the expected benefit from assigning farmers using a targeted policy derived from $\hat{\mu}$. If the same data is used to generate the targeted policy and evaluate its effectiveness, the result will likely overstate the benefits of the policy. We use cross-fitting to address this concern, dividing the data into K equal groups (folds). Denoting $k(i)$ the fold which contains farmer i, for each fold k, we estimate the model on all folds except k, and then use the model parameters to predict pickup for farmers in the k^{th} fold. We repeat this k times to generate $\hat{\mu}_{-k(i)}$ for all farmers in the sample (Figure A6).[13]

The estimate of $\hat{\pi}_{-k}$ also depends on the constraints related to technological limits on the number of farmers that can be called in an hour-day combination. We account for this by using a two-step process. Step 1 uses LASSO and data collected using uniform randomization to predict $\hat{\mu}$ with cross-fitting. Step 2 uses $\hat{\mu}$ and the budget constraints to

[13]The concept of cross-fitting is similar but distinct from cross-validation. Cross-validation is about model selection and selecting tuning parameters where the ultimate goal is to estimate a single model on the entire training data (the one that minimizes the cross-validation error), while with cross-fitting, we retain the models estimated on the different folds as an input to a subsequent step of the analysis. The subsequent step in our context is estimating the targeted policy.

allocate farmers to their optimal or near-optimal call times as shown below:

$$\max_{z} \sum_{i=1}^{n} \sum_{j=1}^{J} z_{ij} \hat{\mu}_{-k(i)}(x_i, j, \cdot) \tag{3}$$
$$s.t. \sum_{j} z_{ij} = 1 \;\; \forall i, \;\; \sum_{i} z_{ij} \leq b_j \;\; \forall j \;\text{ and }\; z_{ij} \in \{0, 1\}$$

where b_j is the capacity limit on call time j. The decision variables are z_{ij}, which takes a value of 1 if farmer i is allocated to call time j and 0 otherwise. The decision variables can be represented in the form of a $N \times J$ matrix z, where each row corresponds to a farmer and each column represents a treatment arm. The set of constraints is combined to generate a matrix A that has dimension $(N + J) \times (N * J)$. The first N rows of A ensure that every farmer i can be allocated to only one call time j. The remaining J rows ensure that the sum of allocation in each call time cannot exceed the total capacity, and, therefore, they should add to be less than or equal to b_j. The resulting set of equations for constraints is $A \times \text{vec}\, z \leq [1, .., 1, b_1, .., b_J]$. Since the decision variables are discrete, this can be set up as an integer programming problem.

We use the Gurobi Parallel Mixed Integer Programming solver. This solver uses a linear-programming-based branch-and-bound algorithm for mixed integer programming problems. All the steps are summarized in the targeting algorithm below for a training dataset $\mathcal{S}$.

Algorithm 1 Algorithm for Estimating $\hat{\pi}$

Input: Capacity limits b; Training data $\mathcal{S}$
Result: Model parameters $(\hat{\delta}, \hat{\beta}, \hat{\gamma})$, Allocation z

1: Partition data into K mutually exclusive folds
2: Estimate LASSO using all folds except k (Equation 2): Obtain $\hat{\delta}, \hat{\beta}, \hat{\gamma}$
3: Predict for each i in k: $\hat{\mu}_{-k}(x_i, J_i, \cdot)$
4: Repeat 1, 2, and 3 for all folds: Append $\hat{\mu}_{(-k)}\ \forall\ k \in K$.
5: Constrained Optimization: $\max_z \sum_{i=1}^{n} \sum_{j=1}^{J} z_{ij} \hat{\mu}_{-k(i)}(x_i, j, \cdot)$ s.t. $\sum_j z_{ij} = 1 \forall i$ and $\sum_i z_{ij} \leq b_j \forall j$ (Equation 3)
6: Use $\hat{\mu}$, capacity limit b and solve for z
7: **return** $[\hat{\delta}, \hat{\beta}, \hat{\gamma}, z]$

Once we have estimated $\hat{\pi}$, we can evaluate it in two ways. In on-policy evaluations, we implement the targeted policy in practice and compare it to the alternative uniform policy. Alternatively, we can employ off-policy or counterfactual evaluation using the data collected via $\mathcal{U}$ to evaluate the counterfactual outcome from implementing any targeted policy. Below we describe the two concepts used to evaluate targeted policies in this project.[14]

6.2 Off-Policy Evaluation

Off-policy evaluation is used to compute the value of targeted policies from a dataset that is collected using a different policy (Athey and Wager, 2021; Zhou et al., 2022). For instance, for this experiment, data is collected using uniform randomization for weeks $t \in \{1, 2, .., 6\}$. This uniform randomized data can be used to estimate and evaluate $\hat{\pi}$.

Appendix D provides details on off-policy evaluation for a simplified setting with two covariates and four treatment arms. The key idea is that under uniform randomization, $\frac{N}{91}$ farmers receive a treatment under $\mathcal{U}$ that matches their assignment under $\hat{\pi}$. The outcome of this subset of farmers can be used to estimate the value of targeted policy counterfactually.

Estimator of Policy Value: The estimators for the population objects of interest (defined in Section 4) are obtained using data collected from the uniform random policy. All targeted policies are based on an underlying outcome model estimated via cross-fitting. The estimator for the value of $\hat{\pi}$ under the off-policy evaluation is denoted by $\hat{V}^{\text{Off}}(\hat{\pi}, \mathcal{S}^{\text{eval}})$.

$$\hat{V}^{\text{Off}}(\hat{\pi}, \mathcal{S}^{\text{eval}}) = \frac{\sum_{i \in \mathcal{S}^{\text{eval}}} Y_{it}^{\text{obs}} \mathbb{1}(\hat{\pi}(x_i, \cdot) = J_{it})}{\sum_{i \in \mathcal{S}^{\text{eval}}} \mathbb{1}(\hat{\pi}(x_i, \cdot) = J_{it})}$$

This is an unbiased estimator of the value of the targeted policy given randomized data and the use of cross-fitting in constructing the policy. Note that after we estimate $\hat{\pi}$ using the uniform randomization data, there is a further important step of conducting the off-policy evaluation, as we want to counterfactually estimate the value of $\hat{\pi}$ prior to deployment. In a production setting, we would only deploy $\hat{\pi}$ if, counterfactually, we see significant gains of deploying $\hat{\pi}$ over our baseline uniform policy.

6.3 On-Policy Evaluation

Once the off-policy evaluation showed that there were significant gains from deploying $\hat{\pi}$ over the uniform policy, the next step was to deploy the targeted policy $\hat{\pi}$. In the subsequent week, we randomly divided the farmers into two groups. Group 1 received calls according to $\hat{\pi}$. Group 2 received calls according to the uniform policy. At the end of the week,

[14]For comparing the value of two policies in the same data set, we cannot analyze the two components of the difference separately since the same observations appear in both terms in regions of overlap. Hence, we sum over observations and for each observation take the difference.

engagement data was collected for both the groups, as illustrated in Figure 1. This data allows us to do on-policy evaluation.

The population object of interest for the on-policy evaluation is defined as $\delta(\hat{\pi},\mathcal{U}) = V(\hat{\pi},\mathcal{S}^{\text{eval}}) - \bar{V}(\mathcal{U},\mathcal{S}^{\mathcal{U}})$, where $\mathcal{S}^{\text{eval}}$ is the evaluation data for the targeted policy and $\mathcal{S}^{\mathcal{U}}$ is the evaluation data for the uniform policy group. In addition to the overall difference between the two policies, these differences can also be computed for the popular hours. In order to do this comparison for call time j, the covariate space is split in a way such that, for the subset of covariates, the best call time is j. This subset and the population objects of interest are defined below. Appendix D explains on-policy evaluation for the simplified setting. We use $R = \{x \in \mathcal{X} : \hat{\pi}(x,\cdot) = j\}$ to denote the subset. Moreover, difference in the value of two policies (δ) is defined below.

$$\delta(\hat{\pi},\mathcal{U})^j = \mathbb{E}_{x\in R}[\mathbb{E}_Y[(Y_t(\hat{\pi}(\cdot))]] - \mathbb{E}_{x\in R}[\mathbb{E}_Y[\mathbb{E}_{J\sim\mathcal{U}}[\sum_j (Y_t(\pi^j)\mathbb{1}(J=j))]]]$$

Note that on-policy estimation is straightforward: it simply entails taking a sample mean on a dataset where the relevant policy was applied. For reference, we define $\hat{V}^{\text{On}}(\mathcal{S}) = \frac{\sum_{i\in\mathcal{S}} Y_{it}^{\text{obs}}}{|\mathcal{S}|}$.

7 Results: Estimating, Evaluating, and Deploying Targeted Policies

In this section, we present our main results employing on- and off-policy evaluation. Our first step was to collect the data using uniform randomization. This uniform randomization data was then used to estimate $\hat{\pi}$. We present the $\hat{\pi}$ that was estimated using the uniform randomization data from weeks 1, 2, and 3 and call it $\hat{\pi}_D$.

The machine learning model used for this analysis is the LASSO model (Equation 2). This model, along with the datasets, were used to predict the probability of pickup for every farmer, for each of the 91 call times ($\hat{\mu}$). We do not penalize the coefficients on the treatment dummies. The optimal penalty (λ_0) for all other coefficients is chosen to minimize cross-validation error. Figure 3 (a) shows the Mean Squared Error (MSE) corresponding to different values of λ.[15]

[15] Panel (a) in Figure 3 shows the output of the cross-validation exercise used to choose the optimal λ. It displays two special values of λ with vertical bars in the graph. The λ corresponding to the vertical bar on the left minimizes the cross-validated error. The λ corresponding to the vertical bar on the right provides the highest penalized model such that the cross-validated MSE is within one standard error of the lowest λ. We use the λ that minimizes the cross-validated error. λ_{1SE} is also reported for extreme cases when too

To assess the out-of-sample prediction performance of the LASSO model, the data can be divided into K folds for cross-fitting. The cross-fitting process is discussed in section 6. Panel (b) in Figure 3 shows the Receiver Operating Characteristics (ROC) plot for the binary pickup (out of sample cross-fitted data). The Area Under the Curve (AUC) is 0.683. This suggests that the LASSO model would be able to correctly predict whether a user will answer the call or not in about 68.3% of the cases. Panel (c) shows the distribution of predicted responses by the class of the binary pickup variable. The predicted response for farmers with actual pickup as 1 is higher when compared to those who did not pick up the calls. Lastly, Panel (d) in Figure 3 shows the calibration plot between the predicted and true pickup. The calibration plots done separately by farmer covariates are shown in Figure A7. We provide the plots by farmer gender and by farmer land size in this appendix.

Following the steps in Algorithm 1 on the uniform randomized data for weeks 1, 2, and 3, off-policy evaluation can be used to estimate the value of $\hat{\pi}_D$ counterfactually. The means and standard errors corresponding to off-policy estimator $\hat{V}^{\mathrm{Off}}(\hat{\pi}_D, \mathcal{S}^{\mathrm{eval}})$ defined in Section 6.2 can be used to estimate the value of targeted policy counterfactually. Moreover, the means and standard errors of the estimator defined in Section 6.3 can be used to estimate the value of the uniform policy. Figure 4 shows that there are substantial gains (2.6 pp or 8%) to calling farmers according to $\hat{\pi}_D$ over the baseline of uniform policy. To provide a sense of the gains for different optimal call times, Figure 4 also reports the off-policy gains for the 12 most popular call times (which comprise 91% of the targeted policy sample.)

Next, we explore whether there are gains from adding additional covariates to the model, incorporating farmer duration, start month, and mean past engagement as covariates.[16] To examine the prediction accuracy of this model, we repeat the same steps as above and compute the ROC curve and AUC. The AUC is 0.681 when we incorporate these additional measures from the historical data. The targeted policy using this model produces gains close to 8%, which is comparable to the gains observed using $\hat{\pi}_D$.

7.1 Deploying and Evaluating $\hat{\pi}$

As part of the experiment, $\hat{\pi}$ was estimated, deployed, and evaluated; we iteratively updated $\hat{\pi}$ as we collected additional uniform data. We also refined some details of our approach. For example, in week 2, we used 21 time slots (morning, afternoon, and evening) over seven days, subsequently shifting to 91 slots (13 hour slots per day), and, in other weeks, holidays

many variables get dropped from the specification. However, we do not have such an extreme situation for our estimation, and we use $\lambda_{\min}$.

[16] The details on the matrix completion exercise are presented in Appendix F.3.

Figure 3: Out-of-Sample Predictive Performance

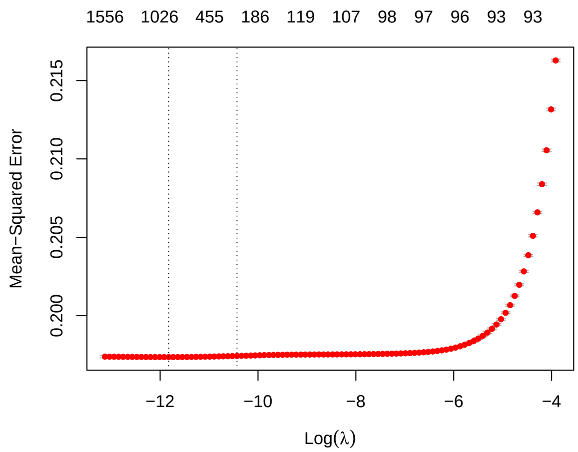

a. Cross-Validation Error

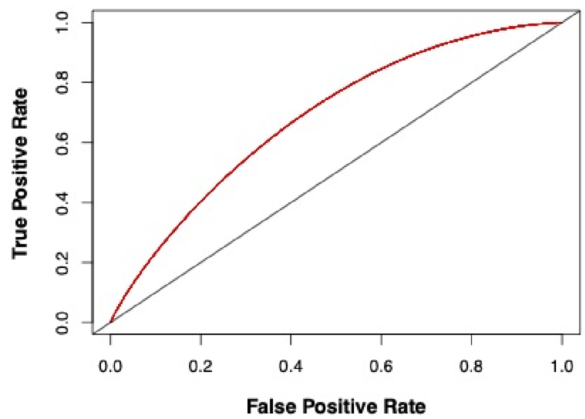

b. ROC curve

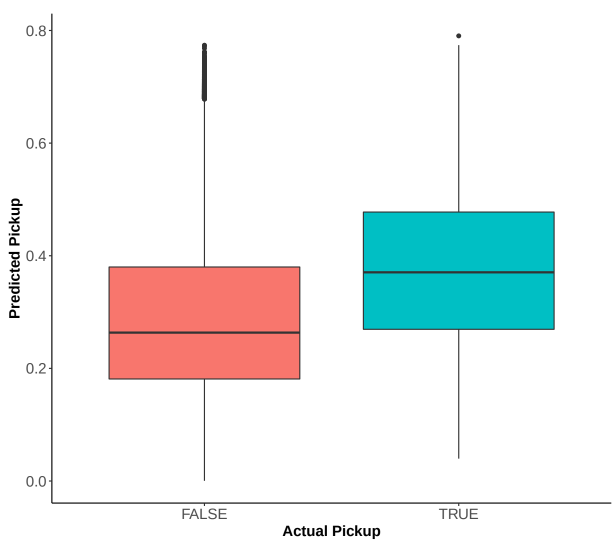

c. Predicted Response

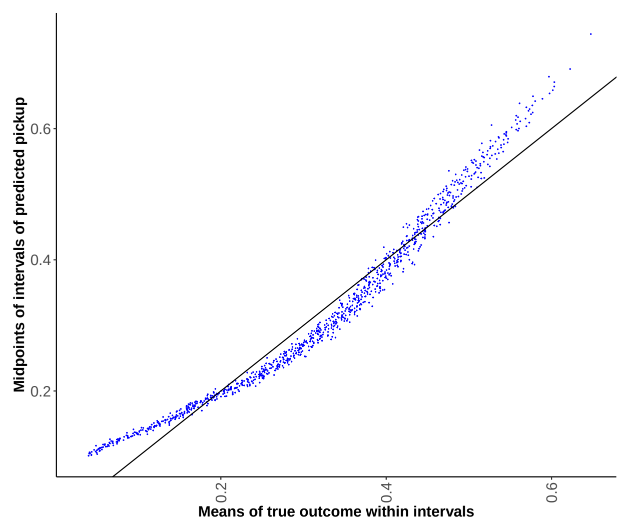

d. Calibration Plot

Notes: (a) shows the MSE corresponding to the different regularization parameters (λ) used in our analysis. (b) shows the ROC curve for the out-of-sample prediction using cross-fitted data. The AUC is 0.683. (c) shows the distribution of predicted response by the binary true pickup in the data. (d) shows the calibration plot for true and predicted pickup using cross-fitted data.

Figure 4: Off-Policy Evaluation: $\hat{\pi_D}$ and Uniform

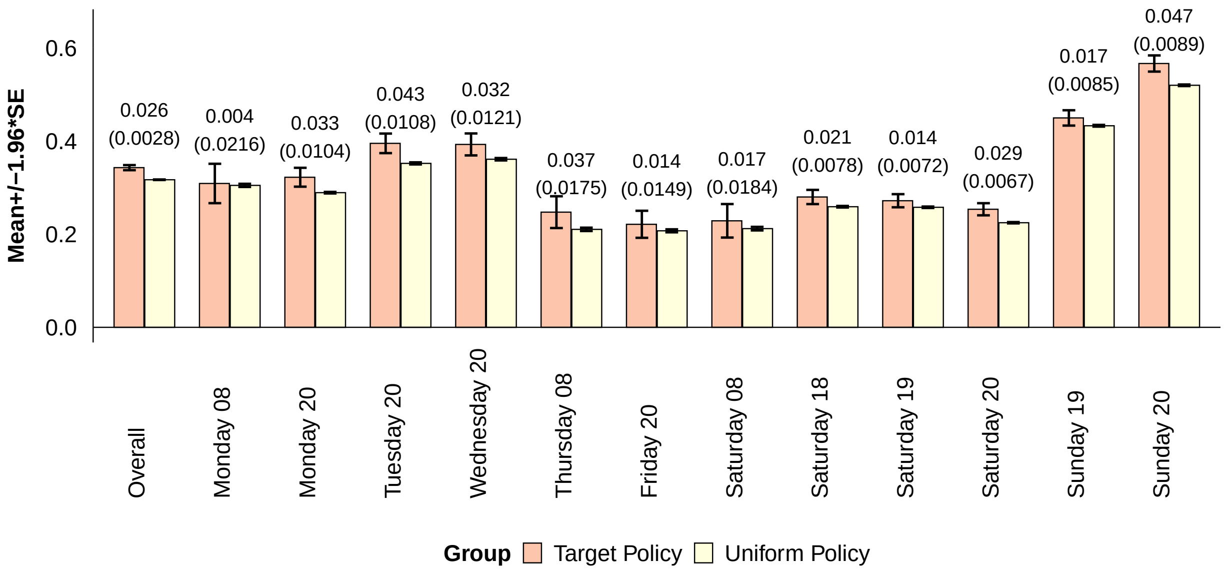

Notes: Sample includes data from uniform randomization in weeks 1, 2, and 3. Targeted policy estimates are computed using $\hat{V}(\hat{\pi}_D, \mathcal{S}^{\text{eval}})$ (Section 6.2). Uniform policy estimates are computed using the sample mean (Section 6.3). The estimate for the differences between $\hat{V}(\hat{\pi}_D, \mathcal{S}^{\text{eval}})$ and $V(\mathcal{U}, \cdot)$ is displayed at the top of the bars and the standard error of this difference is in parentheses.

or other real-world constraints affected the set of days in which we could implement policies. These details are described in Appendix E.

We now present results from the final week of our study, in which we conduct an on-policy evaluation of $\hat{\pi}_D$.[17] In the final week, we randomly assigned farmers into two groups. The first group was assigned targeted policy $\hat{\pi}_D$, and the second group was called according to the uniform randomization. We compare the average pickup for the two groups in Table 3. Table 3 shows the difference between the sample mean of farmers who received calls according to $\hat{\pi}_D$ and those that got called according to the uniform policy. The differences are shown for the overall sample as well for the sample of female and male farmers. Surprisingly, there is only a .4 percentage point increase in overall pickup relative to the control group (30.9 pp); this gain is much smaller than the 2.6 percentage point gains predicted by off-policy evaluation (Figure 4).

Our finding of lower-than-expected performance is in line with the recent work discussing distribution or dataset shifts in the test data (Kuang et al., 2020; Rabanser et al., 2019). In the next two subsections, we seek to understand better why the on-policy per-

[17] Appendix C provides details on the intermediate policies deployed during the experiment in weeks 2, 4, and 5, respectively.

formance was lower than expected and propose approaches to improve the robustness of estimated policies.

Table 3: On-Policy Evaluation in Week 6

Data Collection-policy	$\hat{\pi}_D$	Uniform	Difference
Outcome Variable			
	All Farmers		
First Call	0.313	0.309	0.004
	[0.001]	[0.001]	[0.001]
	Female Farmers		
First Call	0.2971	0.2901	0.0069
	[0.0023]	[0.0014]	[0.0027]
	Male Farmers		
First Call	0.3167	0.3135	0.0032
	[0.0011]	[0.0007]	[0.0013]
N	227,386	587,608	

Notes: On-policy evaluation estimates use sample means for each of the two data collection mechanisms from Week 6.

7.2 Evaluating Targeted Policy on Future Weeks

A robust, targeted policy should work not just out-of-sample for the time period it is estimated but also out-of-sample in subsequent weeks when it is prospectively deployed. However, if the distribution of the outcome variable is subject to systematic variation over time, the policy designer may face a trade-off: placing greater weight on more recent data may focus on behavior that is closer to what will be observed in the near future but at a cost of using less of the information available from earlier weeks.

In this section, we take advantage of the fact that we have a number of weeks of random uniform data, which allows us to vary the number of historic weeks used when estimating a policy and explore the robustness of the policy in subsequent weeks of data (e.g., out-of-sample). This relates to the ideas of stability and transportability of targeted policies on future weeks in (Hitsch et al., 2024). Figure 5 provides an overview of our approach: we consider four scenarios in which older data is down-weighted in our policy estimation. This is possible in a setting such as agriculture where seasonality considerations may affect farmers' workload and task distribution (Gill et al., 1991; Vemireddy and Pingali, 2021). The real world may also be subject to other events, such as festivals or cricket matches, which alter time-use preferences and hence affect actual gains of implemented policies.

Using our notation for training data $\mathcal{S}$ and evaluation data $\mathcal{S}^{\text{eval}}$, we see that the gains from off-policy evaluations do indeed depend on the degree to which the estimation weights more recent vs. less recent data. In Figure 5, "Train" indicates the weeks used to train the policy, and "Test" indicates the week used in the off-policy evaluation. The weights refer to the relative weighting in the LASSO model.[18] The bars give the pickup rate for the targeted policy and the uniform policy. The estimate for the differences in these values is provided at the top of the bars. We find substantial gains (5-8%) for the value of the targeted policy counterfactually estimated over the uniform policy for scenarios where the training data weeks closest to the test week are weighted the highest in the machine learning model. This is likely because the most recent weeks of data capture the technology and preference shocks much better than the older weeks (Figure 5).

Figure 5: Off-Policy Evaluation of Targeted Policies Using Subsequent Weeks' Evaluation Data

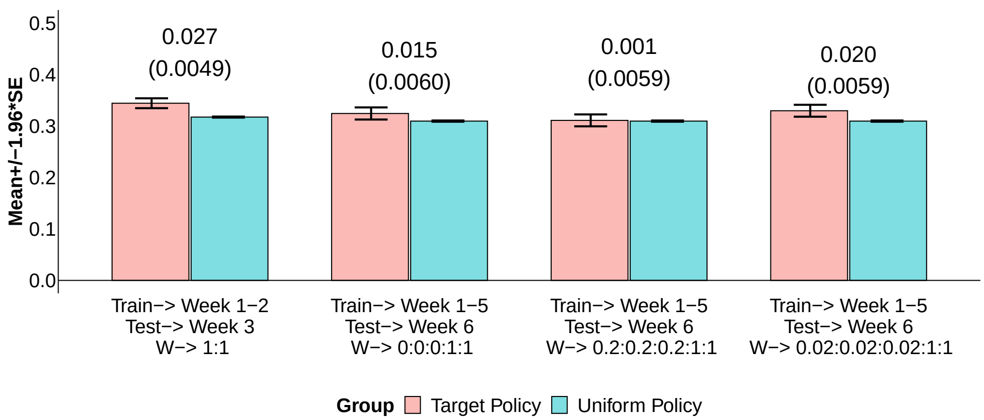

Notes: We use varying subsets of uniform randomization data over the six weeks of the experiment for the estimation ("Train") and evaluation ("Test") of targeted policies in this figure counterfactually. In scenario 1, the targeted policy is estimated with data from weeks 1 and 2, while the evaluation data comes from week 3. The entries for targeted policy in the paper are means and standard errors of $\hat{V}^{\text{Off}}(\hat{\pi}, \mathcal{S}_{\mathcal{U},3})$, as described in Section 6.2. The entries for uniform policy are on-policy sample averages of outcomes for units assigned to uniform policy in the Test week; for scenario 1, this is week 3. Similarly, in scenarios 2, 3, and 4, the targeted policy is estimated using data from weeks 1, 2.., and 5 with varying weights.

[18] The objective function for the weighted LASSO is provided below

$$\sum_t -W_t l(\theta, X_t, J) + \lambda\|\beta\| + \lambda\|\gamma\|$$

where W_t is the weight associated with the training sample in week t and $t \in \{1, 2, .., 5\}$.

7.3 Robustness to Shocks

The overall objective of this project is to maximize farmer engagement with the service. The policy was estimated and deployed for the timing of the first call attempt each week to reach a farmer. In fact, if a farmer does not answer the phone, the dialing software automatically attempts a second call 24 hours later, and if that is not answered, a third and final attempt is made 24 hours later (Figure A8). While, in principle, the second and third calls could have been optimized, the nonprofit did not want to take on such operational complexity. We can, however, use these follow-up calls to provide a sense of how follow-up calls can help mitigate some of the shocks to the first call.

This section explores the efficacy of the rule of thumb of calling farmers 24 hours later and proposes a method for counterfactual evaluation of all three calls. This section illustrates how the targeted policy can improve the overall engagement of farmers relative to the uniform policy. We provide the first evidence for this mechanism using the overall pickup over the three calls for our on-policy evaluation of $\hat{\pi}_D$ (Table 4). We find the impact of targeted policy is higher when examining overall engagement by adding the pickup over call attempts 1, 2, and 3 than only examining the engagement over the first call. If there are shocks to the first call, follow-up calls can mitigate some of the cost.

Table 4: Incorporating Follow-Up Calls to Mitigate Shocks: On-Policy Evaluation

Data Collection-policy	$\hat{\pi}_{\mathrm{D}}$	Uniform	Difference
Outcome Variable			
	All Farmers		
Call 1,2,3	0.552	0.544	0.008
	[0.001]	[0.001]	[0.001]
	Female Farmers		
Call 1,2,3	0.5249	0.5133	0.0116
	[0.0025]	[0.0016]	[0.0029]
	Male Farmers		
Call 1,2,3	0.5577	0.5500	0.0077
	[0.0011]	[0.0007]	[0.0014]
N	227,386	587,608	

Notes: This table shows the on-policy evaluation results for targeted policy $\hat{\pi}_D$ in week 6. The farmers are randomized between two groups. Group A gets called according to $\hat{\pi}_D$, and Group B gets called according to uniform randomization. Sample means are used to estimate the value of $V(\hat{\pi}_D, \mathcal{S}^{\text{eval}})$ and $V(\bar{\mathcal{U}})$.

Next, we provide evidence of the impact of the first-call targeted policy on overall engagement using off-policy evaluations. Figure A9 provides evidence that the benefit of targeting the first calls has benefits for overall farmer engagement. For this exercise, $\hat{\pi}$ is estimated on uniformly randomized data for the first calls in week 1 and week 2. We evaluate the first-call policy not just on the first call but also on the follow-up calls. The sample considered for evaluation for the targeted policy corresponds to a subset of people in the evaluation set whose actual assignment matched the targeted policy according to the first call targeted policy. The first two bars in Figure A9 show the gains of targeting the first call on the first call. Next, we add the pickup over the first and second calls for the evaluation sample. We continue to see that the pickup over the first two calls for the targeted policy group is higher than the pickup for the first and second calls for the uniform group (2 pp (SE=0.36)). Next, we add pickup over calls 1, 2, and 3. We see that the targeted policy engagement is higher than the uniform policy engagement (1.6 pp (SE=0.35)).

Furthermore, we provide some evidence of the benefit of targeting second calls. We estimate a targeted policy for the second call using data on whether farmers answered a second call. Among the farmers who did not pick up the first call in week 1 and week 2, we estimate a second call targeted policy. We evaluate the second call targeted policy on the second call data but adjust for the propensity of the pickup in the first call. Here are the steps to this evaluation: we predict pickup for first calls under the first call counterfactual policy (first call targeted policy). Next, we reweigh all of the second call data in the evaluation step according to the inverse propensity weights. This adjusts for the fact that targeting the first call changes the set of people left over for the second call. The results are provided in Table 5. We observe a 4.3 pp gain from targeting the second calls relative to the baseline mean of the uniform group for the inverse propensity-weighted pickup of 29.8%.

7.4 Bandwidth Constraints and Equity-Efficiency Trade-Off

We begin this section by first estimating the benefits of expanding the bandwidth of the agricultural advisory service. We do this computation counterfactually using the uniform randomization data for weeks 1, 2, and 3 of the experiment. We estimate policies by modifying Algorithm 1, in particular by altering the bandwidth limits in the constrained optimization step. We start with a bandwidth limit of 10,000 farmers per treatment arm and then gradually expand it to 90,000 farmers. We conduct a counterfactual analysis for a scenario that differs from the implemented experiment, where we deployed farmers to

Table 5: Incorporating Follow-Up Calls to Mitigate Shocks

	Targeted Policy	Uniform	Difference
Impact of Policy Targeting Second Call on Pick-up of Second Call			
Second Call	0.341	0.298	0.043
	[0.0061]	[0.0006]	[0.0061]

Notes: The training and evaluation data for estimating the value of targeted policy counterfactually and the uniform policy uses the uniform randomization data for week 1 and week 2. The training data for estimating the second call target policy consists of $N^{\text{second}}_{1,\mathcal{U}} = 630,383, N^{\text{second}}_{2,\mathcal{U}} = 431,290$. Note the estimate for the $\hat{V}^{\text{Off}}(\hat{\pi}, S^{\text{second}}_{\mathcal{U},1,2})$ is done for the inverse propensity weighted second call pick-up instead of the non-adjusted pick-up.

the target and uniform groups. Instead, here we conduct counterfactual analysis for the scenarios where the sample of farmers (600,000 farmers randomly drawn from our final sample in weeks 1,2 and 3 for this exercise) would be called according to the targeted policy, and there is no uniform call time group.

We find substantial counterfactual gains in pickup for the targeted policy by relaxing the budget constraint. As the budget constraint is expanded from 10,000 to 90,000, the estimate for $\hat{V}^{\text{Off}}(\hat{\pi}, \mathcal{S}_{\mathcal{U}})$ increases from 0.3151 (SE=0.0037) to 0.3349 (SE=0.0037). The increase for all farmers is shown in Panel (a) in Figure 6. In Panel (b) in Figure 6, we show the changes for female and male farmers. Finally, in Panel (c), we illustrate the increase separately for smartphone and non-smartphone users. Expanding bandwidth improves call pickup rates for all subgroups. However, since all types of farmers are weighted equally in the optimization, we need large increases in bandwidth to substantially improve the pickup for vulnerable groups such as female farmers and non-smartphone users.

An alternative way to improve the engagement of vulnerable groups with the agricultural advisory service is to attach higher weights to the vulnerable groups in the constrained optimization step. In fact, there is a substantial overlap between the top pickup call times for male and female farmers, as seen in Panel (b) and (c) of Figure 2. This suggests the possibility that an algorithm that is designed to maximize overall pickup rates could assign high-value slots to men; there may be a trade-off in employing a welfare-sensitive algorithm that puts a greater weight on reaching female farmers. This motivates the next analysis, where we vary the weights by farmers' characteristics. We first draw a random sample of 600,000 farmers from the datasets collected using uniform randomization in weeks 1, 2, and 3 of the experiment. We estimate $\hat{\mu}$ as before, modifying the constrained optimization step of Algorithm 1 to add weights as follows. Let G be the group indicator

Figure 6: Relaxing the Budget Constraint

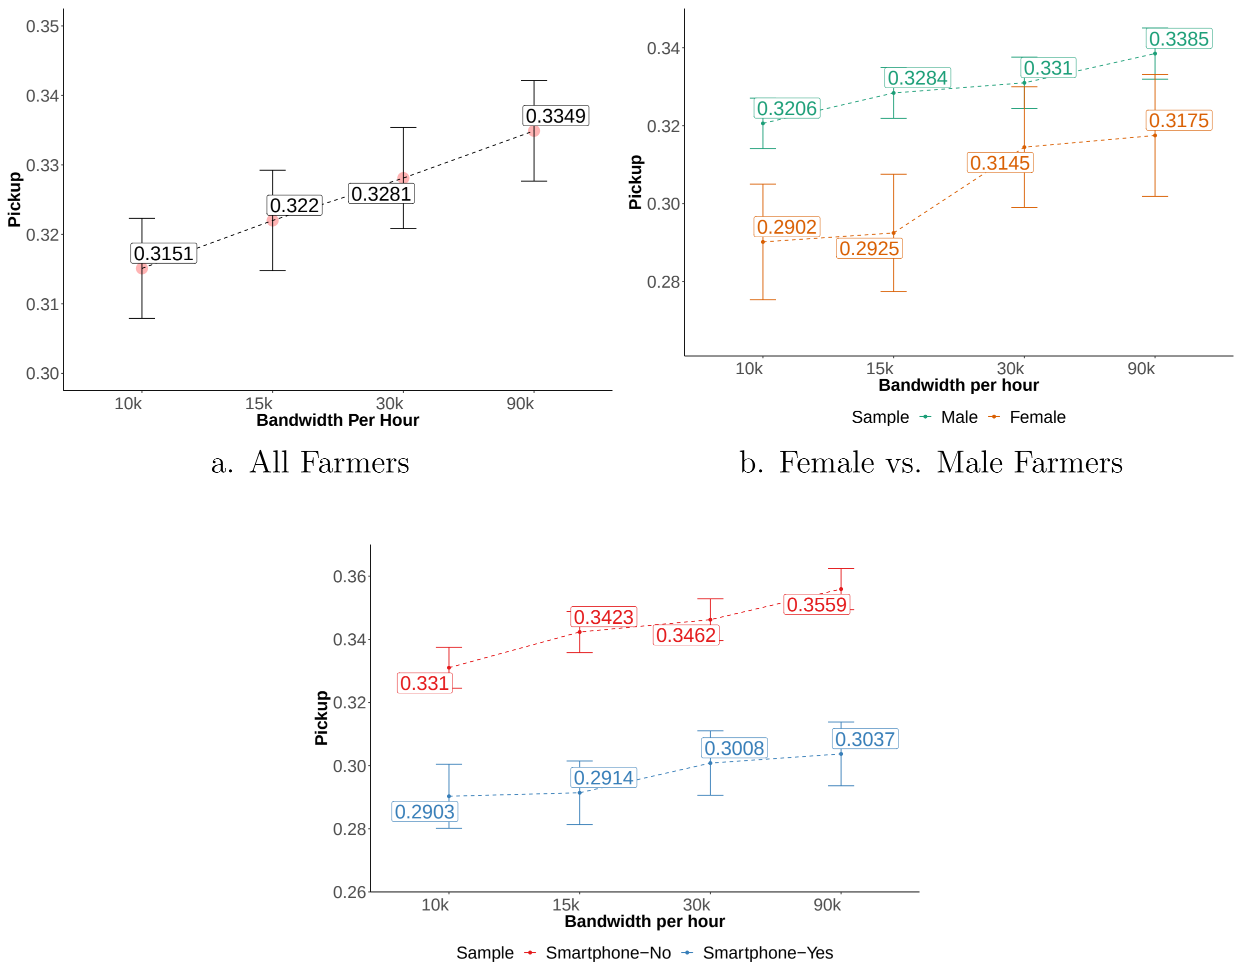

a. All Farmers

b. Female vs. Male Farmers

c. Smartphone vs. Non-Smartphone Users

Notes: The off-policy evaluation is conducted using data collected in the uniform randomization arm during weeks 1, 2 and 3 of the experiment. $\hat{\pi}$ is estimated using Algorithm 1, with modified constraints. Panel (a) shows the gains in the pickup for the sample as the bandwidth is expanded from calling 10,000 farmers in an hour to 90,000 farmers. Panel (b) and (c) show the gains by subgroups.

Figure 7: Weighting Subgroups Differently in Constrained Optimization

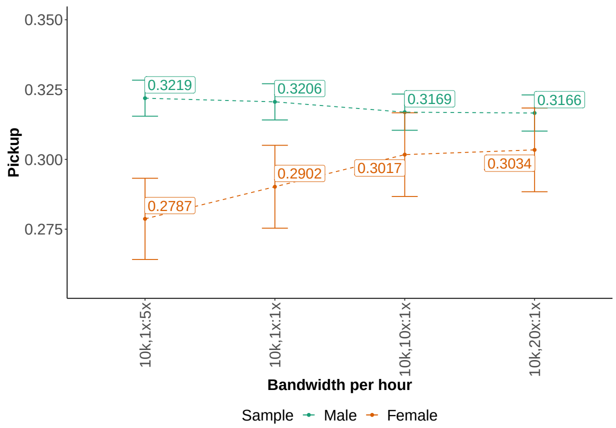

a. Male vs. Female

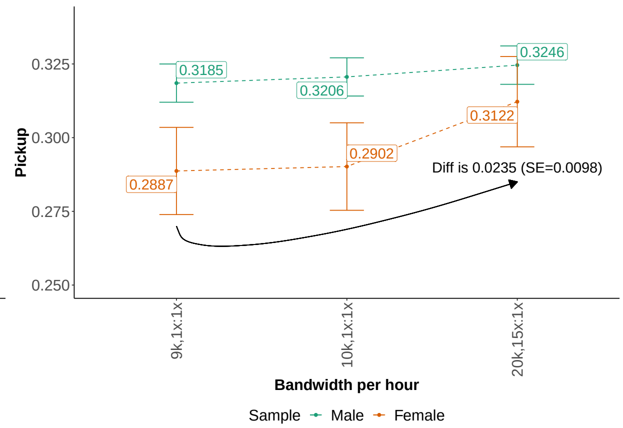

b. Male vs. Female

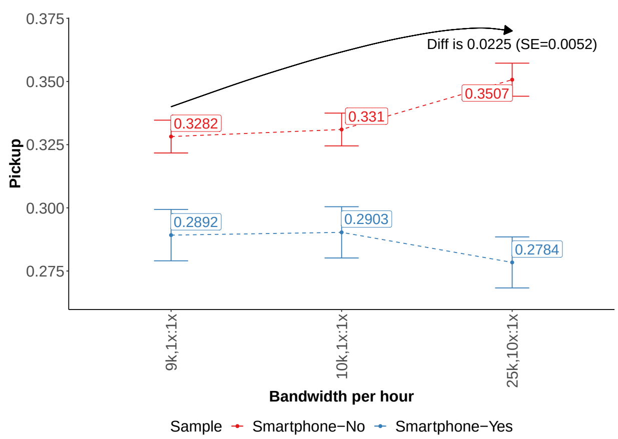

c. Smartphone vs. Non-Smartphone

Notes: The uniform randomization data for weeks 1, 2 and 3 of the experiment are used for these off-policy evaluations. Panel (a) shows the scenarios where the bandwidth is kept constant at 10,000 farmers, and we vary the weights across the three scenarios between male and female farmers. The value of targeted policy is denoted where we attach weights for subgroups is $V(\pi, w, .) = \sum_g w(g) E_{X \sim F_g} E(Y(\pi(X)))$ as defined in Section 7.4. Panel (b) shows an expansion in bandwidth from 9,000 to 20,000 but attaches a higher weight to females than males. Finally, Panel (c) shows gains as we expand the bandwidth but put higher weight on non-smartphone users.

random variable with support in $\mathcal{G}$ and g be the realization of this random variable. X follows a mixture distribution $X|G = g \sim F_g$. The population objective function in this case is $V(\pi, w, \mathcal{S}^{eval}) = \sum_g w(g)\mathbb{E}_{X\sim F_g}[\mathbb{E}_Y Y_t[(\pi(X))]]$, where w is the weight function $w : \mathcal{G} \to [0, 1]$ s.t. $\sum_{g\in\mathcal{G}} w(g) = 1$.

For any given set of weights, we estimate an optimal policy for the corresponding weighted objective function. Panel (a) in Figure 7 illustrates three scenarios for varying weighting schemes for male and female farmers. Here, we keep the bandwidth constant at 10,000 calls per hour and vary weights only. In the first case, males and females are weighted equally. In the second case, females are given 10 times the weight as males. Finally, in the last case, females are given 20 times the weight as males. The constraints are tight and binding in the optimization, which means increasing the relative weight on females moves more female farmers into near-optimal slots. We can also calculate the "real world" sense of trade-offs. Since 18% of our sample is female (sample size is 600,000 for this exercise), moving from scenario 1 to scenario 3 enables us to reach 1,426 more women at the expense of 1,968 fewer men. The bandwidth constraints have been scaled down as the analysis includes a smaller sample than 1.3 million farmers, and we take into account scheduling only one message for each farmer, which also does not include the follow-up calls.[19]

Moreover, we explore the scenarios where the constraints are relaxed, but while doing that, we put higher weights on the vulnerable groups. For instance, in Panel (b) we expand the bandwidth from 9,000 to 20,000 farmers per hour. However, we weight the females 15 times higher than the males in the constrained optimization step. This counterfactual analysis shows that female pickup increases by 2.35 pp (SE=0.98). Similarly, in Panel (c), we expand the bandwidth from 9,000 to 25,000 but weight the non-smartphone users more than the smartphone users. We find that this improves the pickup of non-smartphone users by 2.25 pp (SE=0.52). This translates into reaching an additional 8,505 non-smartphone users at the cost of reaching 2,398 fewer smartphone users. Attaching a higher weight to non-smartphone users as more slots are added could be beneficial as they are likely to benefit more from this advisory service since they might have limited access to other sources of information.

[19] In practice, two or three calls had to be scheduled for 1.3 million farmers each week, and every message call included two follow-up calls if the farmer did not pick up the call on the first and second attempts.

8 Discussion

This project develops, implements, and evaluates a recommendation system designed to increase engagement in an agricultural advisory service in rural areas in an eastern state in India. Our research design, which tests estimated targeted policies against a control group policy of uniform random assignment, allows us to perform both on- and off-policy evaluation. We find that personalizing call timings can increase engagement meaningfully, with estimated off-policy gains as high as 8 percent. These gains come at virtually no programmatic cost: While estimating the targeted policy requires LASSO regression and a complex integer programming problem, the computational cost is negligible, even in a setting with almost one million weekly users.

Our paper serves as an important proof of concept, showing that advanced recommendation systems typically employed in apps or web settings can work with technology as simple as automated voice telephone calls, even in data-poor environments.

Our paper also highlights some of the key trade-offs between on- and off-policy evaluation. Off-policy evaluation has the significant virtue of flexibility. Because our control group was called at randomly assigned times, we were able to tackle a range of important questions. We were able to precisely quantify the equity-efficiency trade-off that the organization would face if it placed higher welfare weights on reaching women farmers, and we show that recent data is more predictive of future farmer behavior than data just a few weeks older.

In contrast, on-policy evaluation has significantly more statistical power. It also helps us to identify the possible existence of technology or preference shocks that degrade the performance of targeted policies in future weeks. We show several ways of modifying the policy to account for the possibility of such shocks. First, attaching a higher weight to training samples closest to the test data week provides robust policies for future weeks. This is because agriculture is a seasonal activity, and the weeks closest to the test data capture the likelihood of such shocks better than older weeks. We also find that the nonprofit's default policy of scheduling follow-up calls 24 hours later and, if necessary, 48 hours later can reduce the apparent effect of shocks. Gains could be even larger if the timing of the follow-up calls were personalized.

Engagement and learning new information are the initial key steps in improving farmers' productivity and welfare. As shown in Fabregas et al. (2019), sharing agricultural information using digital technologies leads to a 22% increase in farmers' odds of adopting recommended agricultural inputs and a 4% increase in farmers' yield. While our experiment focuses on measuring the impact on farmers' engagement with the digital service, future work will evaluate the impact of customizing content on farmers' practices and welfare.

References

Abate, G. T., Abay, K. A., Chamberlin, J., Kassim, Y., Spielman, D. J., and Paul Jr Tabe-Ojong, M. (2023). Digital tools and agricultural market transformation in africa: Why are they not at scale yet, and what will it take to get there? *Food Policy*, 116:102439.

Agrawal, K., Athey, S., Kanodia, A., and Palikot, E. (2022). Personalized recommendations in edtech: Evidence from a randomized controlled trial. *arXiv preprint arXiv:2208.13940*.

Aker, J. C. (2011). Dial "a" for agriculture: A review of information and communication technologies for agricultural extension in developing countries. *Agricultural Economics*, 42(6):631–647.

Aker, J. C., Ghosh, I., and Burrell, J. (2016). The promise (and pitfalls) of ICT for agriculture initiatives. *Agricultural Economics*, 47(S1):35–48.

Anderson, J. R. and Feder, G. (2004). Agricultural extension: Good intentions and hard realities. *The World Bank Research Observer*, 19(1):41–60.

Araya, R., Menezes, P. R., Claro, H. G., Brandt, L. R., Daley, K. L., Quayle, J., Diez-Canseco, F., Peters, T. J., Cruz, D. V., Toyama, M., et al. (2021). Effect of a digital intervention on depressive symptoms in patients with comorbid hypertension or diabetes in Brazil and Peru: Two randomized clinical trials. *JAMA*, 325(18):1852–1862.

Arouna, A., Michler, J. D., Yergo, W. G., and Saito, K. (2020). One size fits all? experimental evidence on the digital delivery of personalized extension advice in nigeria. *American Journal of Agricultural Economics*, 103(2):596–619.

Ascarza, E. (2018). Retention futility: Targeting high-risk customers might be ineffective. *Journal of Marketing Research*, 55(1):80–98.

Athey, S., Bergstrom, K., Hadad, V., Jamison, J. C., Özler, B., Parisotto, L., and Sama, J. D. (2023a). Can personalized digital counseling improve consumer search for modern contraceptive methods? *Science Advances*, 9(40):eadg4420.

Athey, S., Byambadalai, U., Cersosimo, M., Koutout, K., and Nath, S. (2024). The heterogeneous impact of changes in default gift amounts on fundraising. *Available at SSRN 4785704*.

Athey, S., Karlan, D., Palikot, E., and Yuan, Y. (2022). Smiles in profiles: Improving fairness and efficiency using estimates of user preferences in online marketplaces. *arXiv preprint arXiv:2209.01235*.

Athey, S., Keleher, N., and Spiess, J. (2023b). Machine learning who to nudge: causal vs predictive targeting in a field experiment on student financial aid renewal. *arXiv preprint arXiv:2310.08672*.

Athey, S. and Wager, S. (2021). Policy learning with observational data. *Econometrica*, 89(1):133–161.

Baul, T., Karlan, D., Toyama, K., and Vasilaky, K. (2024). Improving smallholder agriculture via video-based group extension. *Journal of Development Economics*, 169:103267.

Beretta, E., Santangelo, A., Lepri, B., Vetrò, A., and Martin, J. C. D. (2019). The invisible power of fairness. how machine learning shapes democracy. In *Canadian Conference on Artificial Intelligence*, pages 238–250. Springer.

Berman, M. and Fenaughty, A. (2005). Technology and managed care: patient benefits of telemedicine in a rural health care network. *Health Economics*, 14(6):559–573.

Blanco, L. R., Hernandez, I., Thames, A. D., Chen, L., and Serido, J. (2023). Mind your money: A community-based digital intervention for improving financial capability among hispanics. *Journal of Economic Behavior & Organization*, 212:629–643.

Bodapati, A. V. (2008). Recommendation systems with purchase data. *Journal of marketing research*, 45(1):77–93.

Bruce, N. I., Murthi, B., and Rao, R. C. (2017). A dynamic model for digital advertising: The effects of creative format, message content, and targeting on engagement. *Journal of marketing research*, 54(2):202–218.

Burlig, F., Jina, A., Kelley, E. M., Lane, G., and Sahai, H. (2024). Long-range forecasts as climate adaptation: Experimental evidence from developing-country agriculture. Working Paper 2024-21.

Casaburi, L., Kremer, M., Mullainathan, S., and Ramrattan, R. (2019). Harnessing ict to increase agricultural production: Evidence from kenya. Working paper. First draft: October 31, 2013; This draft: September 23, 2019.

Cole, S. and Fernando, A. N. (2021). 'Mobile'izing agricultural advice: Technology adoption, diffusion, and sustainability. *Economic Journal*, 131(633):192–219.

Cole, S., Goldberg, J., Harigaya, T., and Zhu, J. (2024). The impact of digital agricultural extension service: Experimental evidence from rice farmers in india. Working paper.

Dammert, A. C., Galdo, J., and Galdo, V. (2013). Digital labor-market intermediation and job expectations: Evidence from a field experiment. *Economics Letters*, 120(1):112–116.

Dammert, A. C., Galdo, J., and Galdo, V. (2015). Integrating mobile phone technologies into labor-market intermediation: A multi-treatment experimental design. *IZA Journal of Labor & Development*, 4(1):1–27.

Devriendt, F., Berrevoets, J., and Verbeke, W. (2021). Why you should stop predicting customer churn and start using uplift models. *Information Sciences*, 548:497–515.

Dubé, J.-P. and Misra, S. (2023). Personalized pricing and consumer welfare. *Journal of Political Economy*, 131(1):131–189.

Dzanku, F. M., Osei, R. D., Nkegbe, P. K., and Osei-Akoto, I. (2020). Information delivery channels and agricultural technology uptake: Experimental evidence from ghana. *European Review of Agricultural Economics*, 49(1):82–120.

Ekeland, A. G., Bowes, A., and Flottorp, S. (2010). Effectiveness of telemedicine: A systematic review of reviews. *International Journal of Medical Informatics*, 79(11):736–771.

Fabregas, R., Harigaya, T., Kremer, M., and Ramrattan, R. (2023). *Digital agricultural extension for development*, pages 187–219. Springer International Publishing.

Fabregas, R., Kremer, M., and Schilbach, F. (2019). Realizing the potential of digital development: The case of agricultural advice. *Science*, 366(6471):eaay3038.

Fernando, A. N. (2021). Seeking the treated: The impact of mobile extension on farmer information exchange in india. *Journal of Development Economics*, 153:102713.

Foster, A. D. and Rosenzweig, M. R. (2010). Microeconomics of technology adoption. *Annual Review of Resource Economics*, 2:395–424.

Friedman, E., Shaddy, F., and Toubia, O. (2023). Understanding fairness perceptions in demographic segmentation and targeting. *Available at SSRN*.

Gill, G. J. et al. (1991). *Seasonality and Agriculture in the Developing World: A problem of the poor and the powerless*. Cambridge University Press.

Giné, X., Townsend, R., and Vickery, J. (2008). Patterns of rainfall insurance participation in rural india. *The World Bank Economic Review*, 22(3):539–566.

Giulivi, N., Harou, A. P., Gautam, S., and Guereña, D. (2023). Getting the message out: Information and communication technologies and agricultural extension. *American Journal of Agricultural Economics*, 105(3):1011–1045.

Haushofer, J., Niehaus, P., Paramo, C., Miguel, E., and Walker, M. W. (2022). Targeting impact versus deprivation. Technical report, National Bureau of Economic Research.

Hitsch, G. J., Misra, S., and Zhang, W. W. (2024). Heterogeneous treatment effects and optimal targeting policy evaluation. *Quantitative Marketing and Economics*, 22(2):115–168.

Hoxby, C. M. (2014). The economics of online postsecondary education: MOOCs, nonselective education, and highly selective education. *American Economic Review*, 104(5):528–33.

Inoue, K., Athey, S., and Tsugawa, Y. (2023). Machine-learning-based high-benefit approach versus conventional high-risk approach in blood pressure management. *International Journal of Epidemiology*, page dyad037.

Jensen, R. (2007). The digital provide: Information (technology), market performance, and welfare in the south indian fisheries sector. *The quarterly journal of economics*, 122(3):879–924.

Knight, H., Jia, R., Ayling, K., Bradbury, K., Baker, K., Chalder, T., Morling, J. R., Durrant, L., Avery, T., Ball, J. K., et al. (2021). Understanding and addressing vaccine hesitancy in the context of COVID-19: Development of a digital intervention. *Public Health*, 201:98–107.

Kuang, K., Xiong, R., Cui, P., Athey, S., and Li, B. (2020). Stable prediction with model misspecification and agnostic distribution shift. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 4485–4492.

Larochelle, C., Alwang, J., Travis, E., Barrerea, V. H., and Andrade, J. M. D. (2019). Did you really get the message? using text reminders to stimulate adoption of agricultural technologies. *Journal of Development Studiess*, 55:548–64.

Lee, J., Foschini, L., Kumar, S., Juusola, J., Liska, J., Mercer, M., Tai, C., Buzzetti, R., Clement, M., Cos, X., et al. (2021). Digital intervention increases influenza vaccination rates for people with diabetes in a decentralized randomized trial. *NPJ digital medicine*, 4(1):1–8.

Li, X., Rong, Y., Zhang, R., and Zheng, H. (2024). Online advertisement allocation under customer choices and algorithmic fairness. *Management Science*.

Liu, C.-W., Wang, W., Gao, G., and Agarwal, R. (2023). The value of virtual engagement: Evidence from a running platform. *Management Science*.

Lu, H., Simester, D., and Zhu, Y. (2023). Optimizing scalable targeted marketing policies with constraints. *Available at SSRN 4668582*.

Ma, W., McKay, A., Rahut, D. B., and Sonobe, T. (2023). An introduction to rural and agricultural development in the digital age. *Review of Development Economics*, 27(3):1273–1286.

Magruder, J. R. (2018). An assessment of experimental evidence on agricultural technology adoption in developing countries. *Annual Review of Resource Economics*, 10:299–316.

Mazumder, R., Hastie, T., and Tibshirani, R. (2010). Spectral regularization algorithms for learning large incomplete matrices. *The Journal of Machine Learning Research*, 11:2287–2322.

Mitra, S., Mookherjee, D., Torero, M., and Visaria, S. (2018). Asymmetric information and middleman margins: An experiment with indian potato farmers. *Review of Economics and Statistics*, 100(1):1–13.

Neslin, S. A., Gupta, S., Kamakura, W., Lu, J., and Mason, C. H. (2006). Defection detection: Measuring and understanding the predictive accuracy of customer churn models. *Journal of marketing research*, 43(2):204–211.

Olaya, D., Vásquez, J., Maldonado, S., Miranda, J., and Verbeke, W. (2020). Uplift modeling for preventing student dropout in higher education. *Decision support systems*, 134:113320.

Owusu, V., Donkor, E., and Owusu-Sekyere, E. (2018). Accounting for the gender technology gap amongst smallholder rice farmers in northern Ghana. *Journal of Agricultural Economics*, 69(2):439–457.

Quisumbing, A. R. and Pandolfelli, L. (2010). Promising approaches to address the needs of poor female farmers: Resources, constraints, and interventions. *World Development*, 38(4):581–592.

Rabanser, S., Günnemann, S., and Lipton, Z. (2019). Failing loudly: An empirical study of methods for detecting dataset shift. *Advances in Neural Information Processing Systems*, 32.

Rafieian, O. and Yoganarasimhan, H. (2021). Targeting and privacy in mobile advertising. *Marketing Science*, 40(2):193–218.

Rambachan, A., Kleinberg, J., Ludwig, J., and Mullainathan, S. (2020). An economic perspective on algorithmic fairness. In *AEA Papers and Proceedings*, volume 110, pages 91–95.

Rodriguez-Segura, D. (2022). Edtech in developing countries: A review of the evidence. *The World Bank Research Observer*, 37(2):171–203.

Sadish, D., Adhvaryu, A., and Nyshadham, A. (2021). (mis) information and anxiety: Evidence from a randomized covid-19 information campaign. *Journal of Development Economics*, 152:102699.

Simester, D., Timoshenko, A., and Zoumpoulis, S. I. (2020a). Efficiently evaluating targeting policies: Improving on champion vs. challenger experiments. *Management Science*, 66(8):3412–3424.

Simester, D., Timoshenko, A., and Zoumpoulis, S. I. (2020b). Targeting Prospective Customers: Robustness of Machine-Learning Methods to Typical Data Challenges. *Management Science*, 66(6):2495–2522.

Spielman, D., Lecoutere, E., Makhija, S., and Van Campenhout, B. (2021). Information and communications technology (ict) and agricultural extension in developing countries. *Annual Review of Resource Economics*, 13(1):177–201.

Swanson, B. E. and Davis, K. (2014). Status of agricultural extension and rural advisory services worldwide. *GFRAS Summary Report*.

Takahashi, K., Muraoka, R., and Otsuka, K. (2020). Technology adoption, impact, and extension in developing countries' agriculture: A review of the recent literature. *Agricultural Economics*, 51:31–45.

Vemireddy, V. and Pingali, P. L. (2021). Seasonal time trade-offs and nutrition outcomes for women in agriculture: Evidence from rural India. *Food policy*, 101:102074.

Voss, C., Schwartz, J., Daniels, J., Kline, A., Haber, N., Washington, P., Tariq, Q., Robinson, T. N., Desai, M., Phillips, J. M., et al. (2019). Effect of wearable digital intervention for improving socialization in children with autism spectrum disorder: A randomized clinical trial. *JAMA pediatrics*, 173(5):446–454.

Yang, J., Eckles, D., Dhillon, P., and Aral, S. (2023). Targeting for long-term outcomes. *Management Science*.

Yardley, L., Spring, B. J., Riper, H., Morrison, L. G., Crane, D. H., Curtis, K., Merchant, G. C., Naughton, F., and Blandford, A. (2016). Understanding and promoting effective engagement with digital behavior change interventions. *American journal of preventive medicine*, 51(5):833–842.

Yoganarasimhan, H., Barzegary, E., and Pani, A. (2023). Design and evaluation of optimal free trials. *Management Science*, 69(6):3220–3240.

Zhou, B. and Zou, T. (2023). Competing for recommendations: The strategic impact of personalized product recommendations in online marketplaces. *Marketing Science*, 42(2):360–376.

Zhou, Z., Athey, S., and Wager, S. (2022). Offline multi-action policy learning: Generalization and optimization. *Operations Research*.

Appendix

A Example Script of Agricultural Advisory Messages

This section provides two examples of agricultural advisory messages that were sent through the extension service to the farmers in our multistage experiment sample.

Example 1. Advisory on pest management: *Namaskar. Today we will discuss about neck blast disease and its management in paddy crop. Due to high relative humidity and differential day and night temperature Neck Blast disease incidence can be seen in paddy crops. To manage these diseases, first drain out excess water from the paddy field. Spray Hexaconazole (Contaf Plus/Hexadhan Plus/Trigger Pro) @ 400-ml/acre or Azoxystrobi+ Difenoconazole (Amistar Top/ Chemistar /Karishma) @ 200-ml/acre or Tebuconazole + Trifloxystrobin (Nativo) @ 80-gram/acre. Thank you and remember that if you have questions about this advisory or need more information, you can call the hotline number on 155333.*

Example 2. Advisory on basal fertilizer application for transplanting: *Namaskar. Today we will share a few tips for applying basal fertilizers correctly for farmers who are growing HYV paddy. If you have not yet done so, we advise you to complete your transplanting by August 15th. Before applying fertilizers at the time of sowing, you should determine what kind of soil you have. This is because the fertilizers recommended are different for different soil. You should apply 35 kg DAP, 30 kg Potash, and 8 kg Urea per acre during the last puddling. Remember again that you should apply 35 kg DAP, 30 kg Potash, and 8 kg Urea per acre at the time of last puddling. Please remember 1acre=25 guntha. However, you should apply Potash in two equal splits at the basal and PI stages if you have sandy soil. Also, do not forget that zinc deficiency is the most widespread micronutrient disorder in paddy, affecting plant growth. For this reason, we suggest that you can apply 10 kg of Zinc Sulphate micronutrient per acre based on a soil test report or if your soil is deficient in zinc. Thank you and remember that if you have questions about this advisory or need more information, you can call the hotline number on 155333, and we will provide this message to your mobile via SMS.*

B Additional Tables and Figures

Figure A1: Number of push calls sent during each call hour in the past

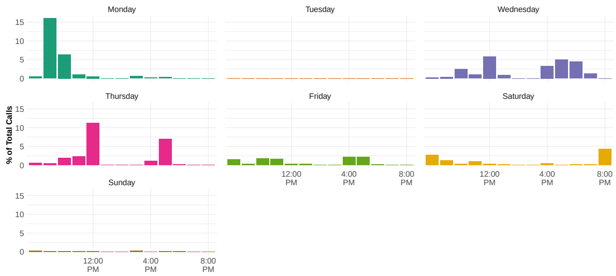

a. 2018

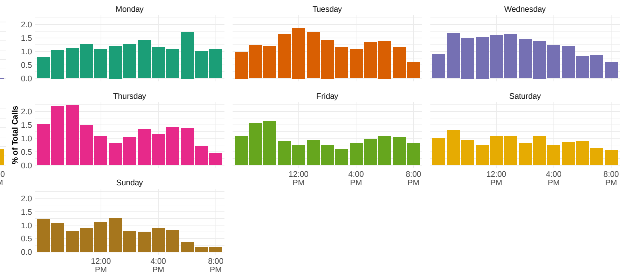

b. 2019

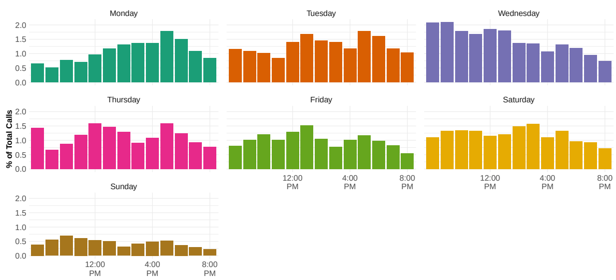

c. 2020

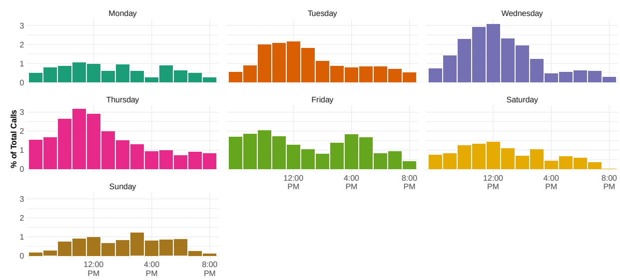

d. 2021

Notes: These figures show the number of push calls sent between July 2018 and September 30, 2021. It illustrates that push calls used to be sent in an ad hoc manner in the past.

Figure A2: Pickup by 91 Treatment Arms

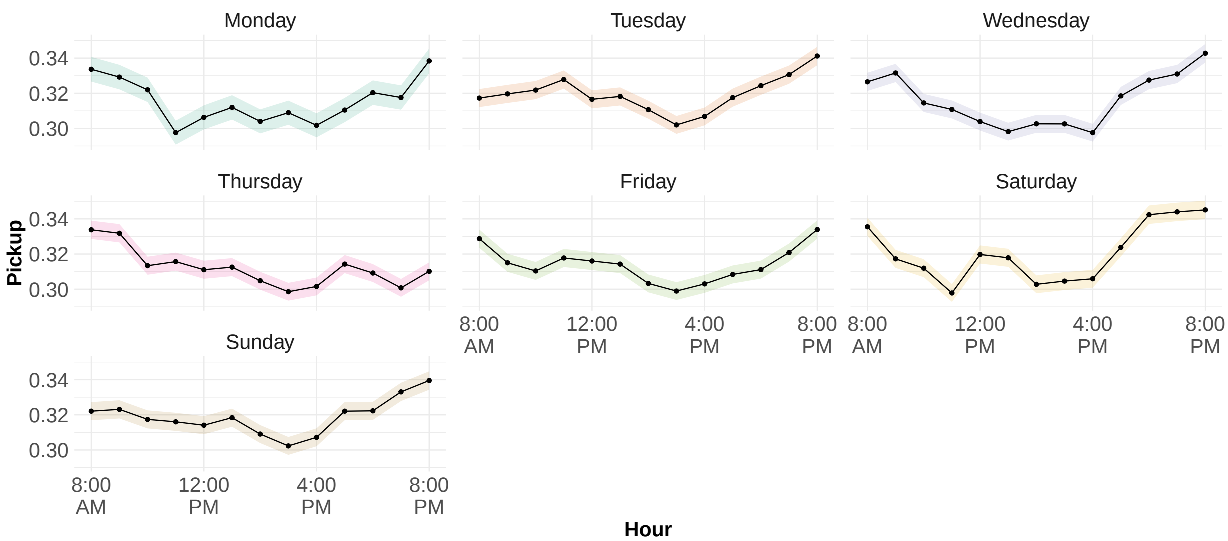

Notes: We pool the data collected using uniform randomization in weeks 1, 2, and 3 of the experiment (see Table 1 for sample size). It illustrates the variation in engagement between the 91 call times, which is a combination of the day of the week and the hour of the day (Mean +/- 1.96SE).

Figure A3: (a) Pickup by time and smartphone, (b) Pickup by the distribution of land size.

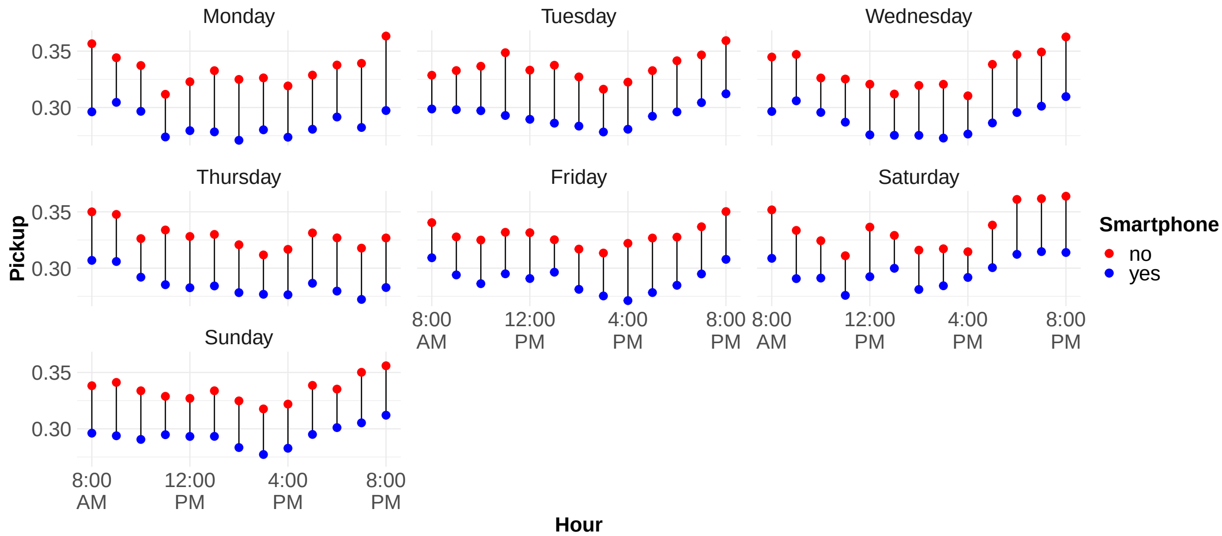

a. Variation in Pickup by Smartphone Dummy for 91 Treatment Arms

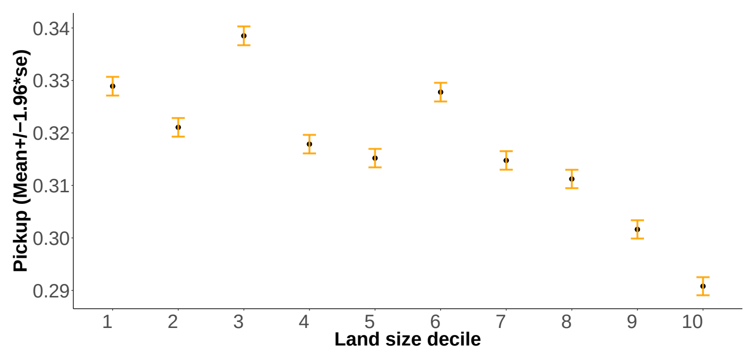

b. Pickup by Land size

Notes: For this figure, we pool the data collected using uniform randomization in weeks 1, 2, and 3 of the experiment. Panel (a) illustrates the smartphone users' and nonusers' gap in engagement with the service over the 91 call times. Panel (b) shows the variation in pickup by land size (Mean +/- 1.96se). We divide the sample of farmers into deciles based on their land size. For each decile, we compute the mean and standard error for pickup.

Table A1: Summary Statistics

Variable	Mean	Std.	Median
	A. Outcome Variable		
Pick-up	0.318	0.283	0.250
	B. Covariates		
Female	0.181	0.385	0.000
Smartphone	0.374	0.484	0.000
Irrigation	0.444	0.497	0.000
Landsize	1.286	1.881	0.809
N	884,194		

Notes: We also include the district of residence and historical engagement as covariates for estimating $\hat{\pi}(x_i, \cdot)$.

Figure A4: Pickup by 91 Treatment Arms

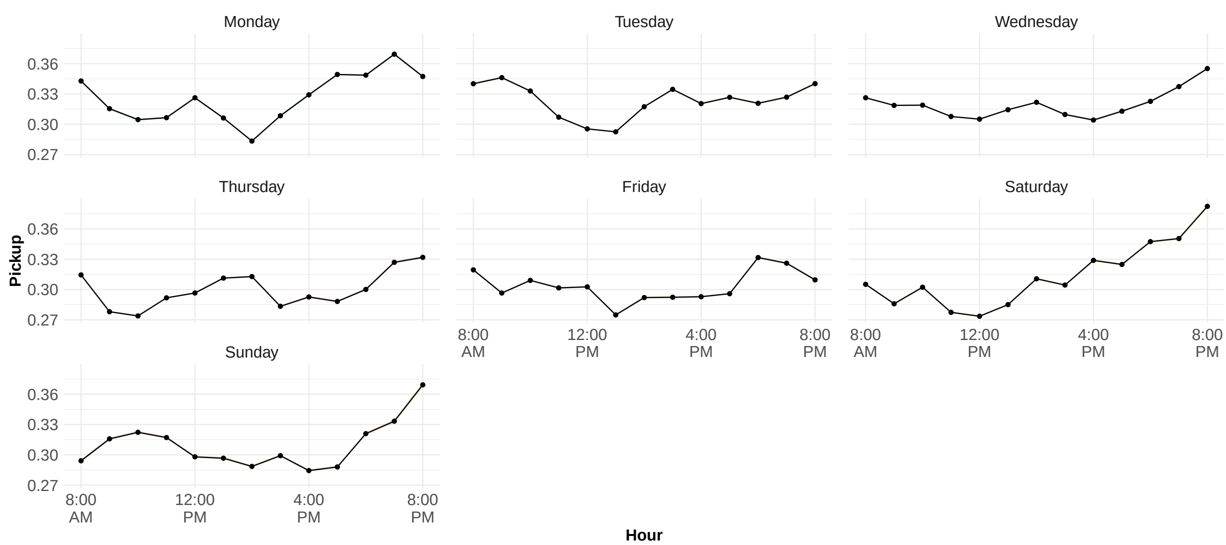

Notes: We pool the data from July 2018 to September 2021 for this figure.

Figure A5: (a) Duration with service, (b) Start month of service, (c) Pickup by time and tenure.

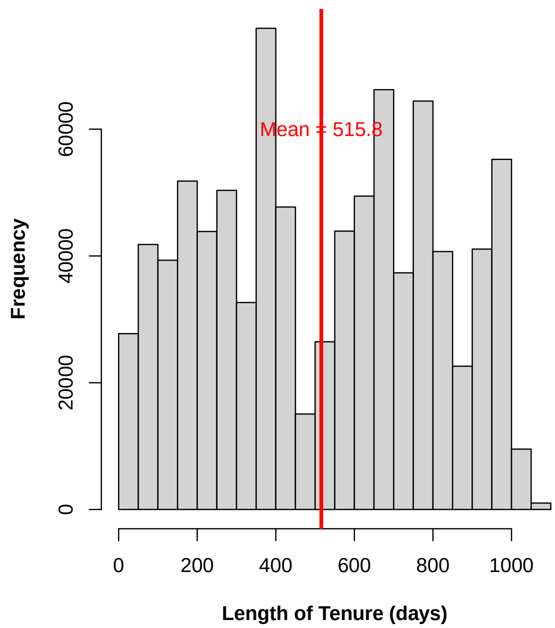

a. Duration (days)

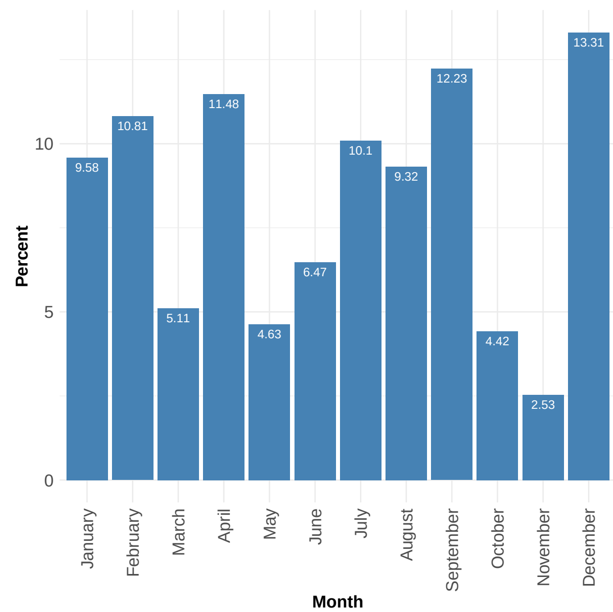

b. Start Month with Service

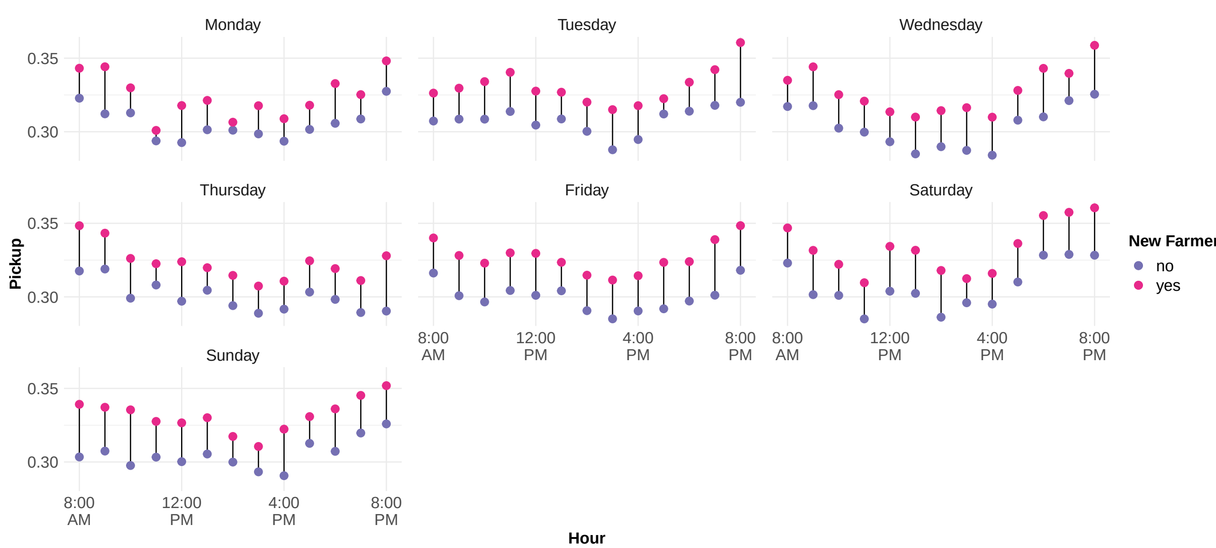

c. Variation in Pickup by New Farmer Dummy for 91 Treatment Arms

Notes: Panel (a) illustrates the distribution of the length of the farmers' tenure with the service. Panel (b) shows the starting month with the service for the sample. For Panel (c), we split the sample into two groups based on their length of tenure. The new farmer dummy takes a value of 1 if the length is lower or equal to the mean length of 516 days and 0 otherwise.

Figure A6: Process of Cross-Fitting on Data (No. of Folds=4).

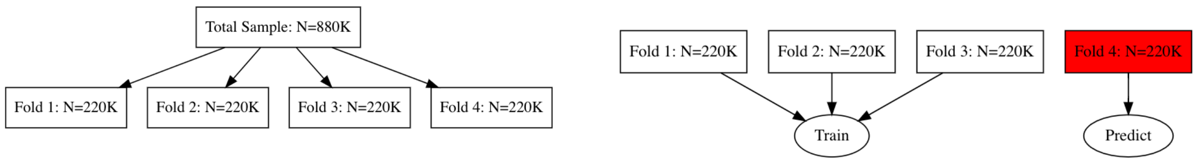

Notes: The figure on the left shows the process of splitting the sample into four folds. The figure on the right shows that the same data is not used for estimation and evaluation. To predict the pickup for the farmers in fold four, the engagement data for farmers in the remaining three folds are used. The same step is repeated for the other three folds.

Figure A7: Calibration Plots for LASSO by Farmer Covariates

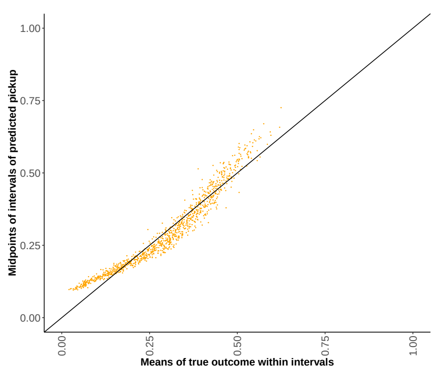

a. Female Farmers

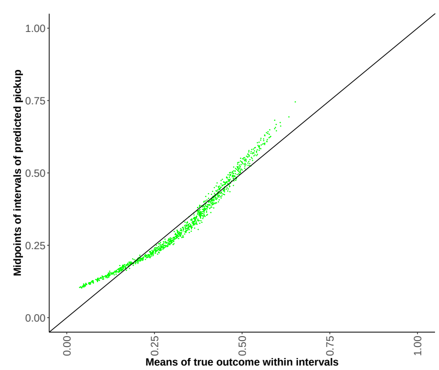

b. Male Farmers

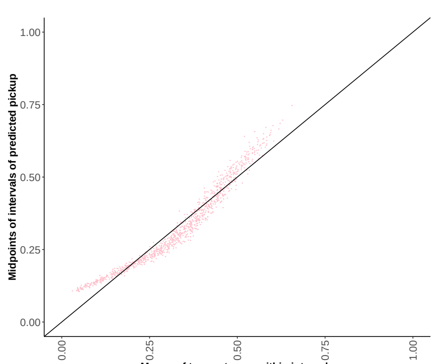

c. Small Landowners

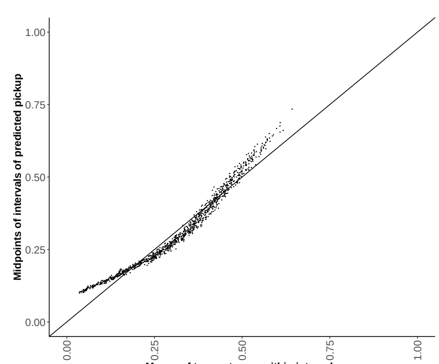

d. Large Landowners

Notes: Panel (a) shows the calibration plots between the true and predicted outcome for the female farmers. Panel (b) shows the calibration plots for the male farmers. Panel (c) shows the calibration plot for the small landowners whose land size is below or the same as the 25^{th} percentile of the land-size distribution. Panel (d) shows the calibration plot for large landowners. Large landowners are defined as those farmers whose land holding is above 25^{th} percentile of the land-size distribution.

Figure A8: Repeat Calls from the Call Center

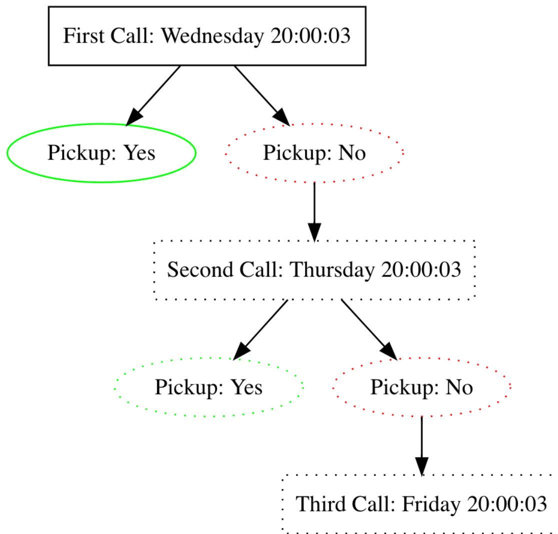

Notes: This figure illustrates the process of follow-up calls from the call center. If the farmer does not pick up the call on the first attempt, the call center makes a call exactly 24 hours after the first call. Moreover, if the farmer does not pick up the call on the second attempt, the third call is made 24 hours after the second call.

Figure A9: Impact of Targeting First Call on Overall Farmer Engagement: Off-Policy Evaluation

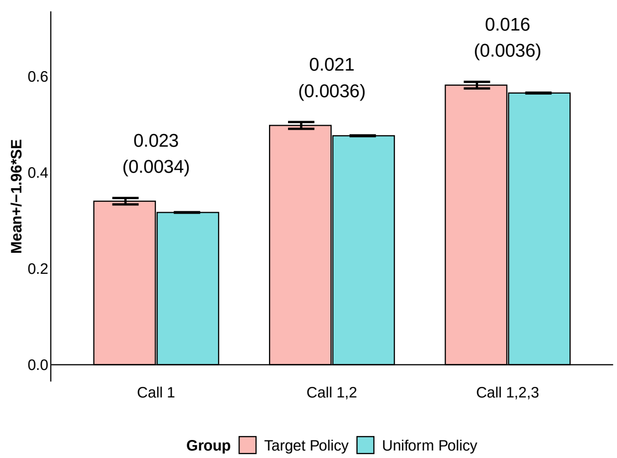

Notes: The training and evaluation data for estimating the value of $\hat{\pi}$ counterfactually and the uniform policy uses the uniform randomization data for week 1 and week 2. We estimate $\hat{\pi}$ on the first calls for the two weeks of data using cross-fitting (following Algorithm 1) and counterfactually evaluate $\hat{\pi}$ using the estimator defined in Section 6.2. We estimate $\hat{V}^{\text{Off}}(\hat{\pi}, \mathcal{S}_{\mathcal{U},1,2})$, where the outcome changes to the total pickup for call 1, call 2 for second pair of bars. The evaluation dataset for the $\hat{\pi}$ group corresponds to the subset of people whose actual assignment matched the first call $\hat{\pi}$. We estimate $\hat{V}^{\text{Off}}(\hat{\pi}, \mathcal{S}_{\mathcal{U},1,2,3})$, where the outcome changes to the overall pickup for call 1, 2, 3 for right most pair of bars.

C Power Calculations

This section provides an outline of the power calculations that we did before starting the multistage experiment. The goal of this exercise is to assess the minimum detectable effect for the off-policy and on-policy evaluations that we intended to do for each week. For each N, we estimate the minimum detectable effect for the differences in the estimate of the value of targeted policy and uniform policy under off-policy and on-policy evaluations. For the off-policy evaluation, let us assume that 900,000 farmers are allocated to uniform randomization in a week. This implies about $1/91 \times 900,000$ would be receiving calls as per the targeted policy under the uniform randomization data (refer to simplified setting in Section 6.2 for details on the fraction). Varying the sample size for those getting called according to $\hat{\pi}$ and those according to the uniform policy, we illustrate the relationship between the minimum detectable effect and the sample size. Figure A10 shows that under the sample sizes close to the ones we see in our project, the MDE for off-policy evaluation is much higher than the MDE for on-policy evaluation. This is happening as we can allocate a higher sample size (250,000) to receive calls according to the targeted policy under on-policy evaluation.

Figure A10: Power Calculations

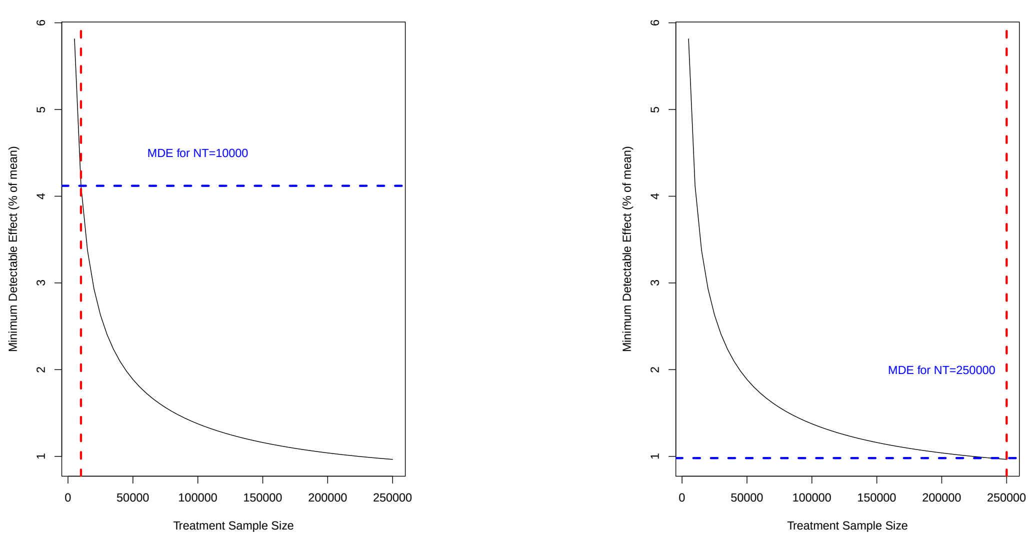

a. Off-Policy Evaluation

b. On-Policy Evaluation

Notes: Panel (a) shows the relationship between the MDE and sample size for the off-policy evaluation. Panel (b) shows the relationship between MDE and sample size for the on-policy evaluations.

D Simplified Setting

D.1 Off-Policy Evaluation

The concept of off-policy evaluation can be explained using a simplified setting with four treatment arms and two covariates. Assume two uniformly distributed covariates $X = [X_1, X_2]$ and the outcome variable is Y. Let there be only four treatments (W_1, W_2, W_3, W_4). The two policies $\mathcal{U}$ and $\hat{\pi}$ are defined below.

$$\mathcal{U} = \begin{cases} W_1, & p = 0.25 \\ W_2, & p = 0.25 \\ W_3, & p = 0.25 \\ W_4, & p = 0.25 \end{cases} \qquad \hat{\pi} = \begin{cases} W_1, & x_1 > 0.5, x_2 > 0.5 \\ W_2, & x_1 > 0.5, x_2 < 0.5 \\ W_3, & x_1 < 0.5, x_2 < 0.5 \\ W_4, & \text{otherwise} \end{cases}$$

Under the uniform randomization, every individual could be allocated to any of those four treatments with a 0.25 probability. Under the targeted policy, individuals in the top left quadrant, top right quadrant, bottom left quadrant, and bottom right quadrant would be assigned to treatment W_4, W_1, W_3, and W_2, respectively (Figure A11). Because of uniform randomness, $\frac{N}{\text{No. of arms}}$ individuals under the uniform randomization are actually assigned to the treatment arm they would have received if the targeted policy $\hat{\pi}$ was deployed. Those individuals are highlighted using blue points in the figure, and they would be the treated individuals in the off-policy evaluation.

Figure A11: Off-Policy Evaluation: Simplified Setting

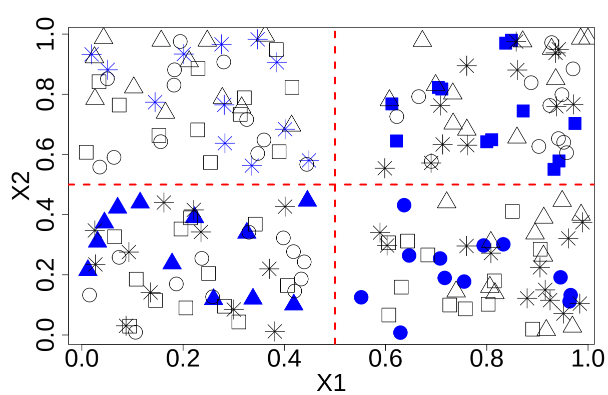

Notes: Data is collected using the uniform randomization, where probability of allocation to each arm is 0.25. 25% of farmers get the same assignment under $\mathcal{U}$ as they would have received under $\hat{\pi}$ (blue points).

D.2 On-Policy Evaluation

Using the same simplified setting with two covariates, we illustrate graphically the concepts of on-policy evaluation used in this study. Figure A12 illustrates the policy for the above setting. The difference in expectation for the subset of individuals who should get W_1 based on $\hat{\pi}$ is obtained by comparing the outcomes for the upper right quadrant in the uniform policy group and the targeted policy group. In other words, the covariate space is kept the same and only the policy is varied between the two groups. Similarly, the differences can be computed for the other arms. Additionally, the overall benefit of deploying $\hat{\pi}$ over uniform policy can be computed by comparing outcomes in Panel (a) with the outcomes in Panel (b).

Figure A12: On-Policy Evaluation: Simplified Setting

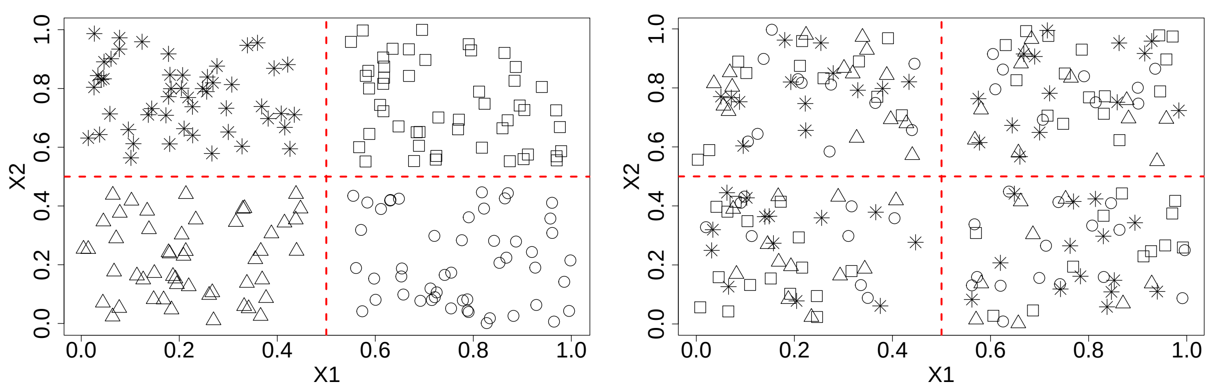

a. Deploy $\hat{\pi}$ b. Deploy $\mathcal{U}$

Notes: Here, individuals are randomly allocated between the group that receives the treatment according to $\hat{\pi}$ (left) and uniform randomization (right).

E Estimated Policies: Implementation Phase

In this section, we provide details on the implementation phase of the targeted policies. While collecting the data using the uniform randomization, we simultaneously developed the technology for deployment of targeted policies. The goal of this exercise was to test the technological constraints with the advisory system and update the targeted policies according to our learning. For the second week, we estimated and evaluated a crude policy ($\hat{\pi}_A$), where the treatment arms were morning, afternoon, and evening trained on uniform randomization data in week 1 (refer to Figure 1).

Next, as we collected more data using the randomization between 91 call times for additional weeks, we updated the policy ($\hat{\pi}_B$) to take into account the hour of the call along with the day of the week. We could only test the updated policy for a limited number of days for the week of November 1st to 7th (week 4) because, during this week, the call center only operated for five days instead of seven (this week had two major festivals, Diwali and Bhai Dooj). Next, we tested another intermediate policy ($\hat{\pi}_C$) in week 5. This was the first week that we could implement the targeted policy with 91 treatment arms for all seven days. However, during this time, we were also trying to learn about the capacity constraints and bandwidth limits for the call center. We deployed and tested the intermediate policy for week 5 but with tighter constraints. For the last week, week 6, we deployed the policy for all seven days and relaxed the constraint further to accommodate additional farmers in their best-predicted call times ($\hat{\pi}_D$). Ideally, we would want to implement a few more targeted policies taking into account the possibility of shocks, but we used off-policy evaluation to estimate and evaluate those additional policies. We only had time to deploy a few limited, targeted policies during the six weeks.

Table A2: Implementation Phase (Week 2)

Data Collection Method	$\hat{\pi}_A$	Uniform Policy	Difference
All	0.3030	0.3178	-0.0148
	[0.0009]	[0.0006]	[0.0011]
N	265,188	616,656	

Notes: This table shows the on-policy evaluation for $\hat{\pi}_A$ in week 2. The farmers are randomized between two groups. Group A gets called according to $\hat{\pi}_A$ and group B gets called according to uniform randomization. Sample means are used to estimate the value of $V(\hat{\pi_A}, \mathcal{S}^{\text{eval}})$ and $\bar{V}(\mathcal{U}, \mathcal{S}^{\mathcal{U}})$.

Table A3: Implementation Phase (Week 4)

Data Collection Method	$\hat{\pi}_B$	Uniform Policy	Difference
All	0.3125	0.3172	-0.0047
	[0.0011]	[0.0006]	[0.0013]
N	167,995	707,644	

Notes: This table shows the on-policy evaluation for $\hat{\pi}_B$ in week 4. The farmers are randomized between two groups. Group A gets called according to $\hat{\pi}_B$ and Group B gets called according to uniform randomization. Sample means are used to estimate the value of $V(\hat{\pi_B}, \mathcal{S}^{\text{eval}})$ and $\bar{V}(\mathcal{U}, \mathcal{S}^{\mathcal{U}})$.

Table A4: Implementation Phase (Week 5)

Data Collection Method	$\hat{\pi}_C$	Uniform Policy	Difference
All	0.3133	0.3195	-0.0063
	[0.0011]	[0.0006]	[0.0011]
N	234,276	624,801	

Notes: This table shows the on-policy evaluation for $\hat{\pi}_3$ in week 5. The farmers are randomized between two groups. Group A gets called according to $\hat{\pi}_C$ and Group B gets called according to uniform randomization. Sample means are used to estimate the value of $V(\hat{\pi_C}, \mathcal{S}^{\text{eval}})$ and $\bar{V}(\mathcal{U}, \mathcal{S}^{\mathcal{U}})$.

F Alternative Specifications

F.1 Functional Form of the Treatment Variables

For the main analysis, the treatment effects are incorporated as 91 treatment dummies in our specifications. We estimate an alternative specification here. We define the hour variable as a continuous variable and use a polynomial function. There are also seven dummies, one for each day of the week. The model incorporates interactions of the hour variable and its higher-order terms with the days of the week. The modified model is provided below.

$$\text{logit}(p_{ij}) = \beta_1 X_i + \beta_2 \text{hour}_j + \beta_3 \text{hour}_j^2 + \beta_4 \text{hour}_j^3 + \beta_5 \text{day}_j + \eta_1 \text{hour}_j \text{day}_j + \eta_2 \text{hour}_j^2 \text{day}_j + \eta_3 \text{hour}_j^3 \text{day}_j + \gamma_1 X_i \text{hour}_j + \gamma_2 X_i \text{hour}_j^2 + \gamma_3 X_i \text{hour}_j^3 + \gamma_4 X_i \text{day}_j + \delta_1 X_i \text{hour}_j \text{day}_j + \delta_2 X_i \text{hour}_j^2 \text{day}_j + \delta_3 X_i \text{hour}_j^3 \text{day}_j \tag{4}$$

The parameters of the model are estimated using LASSO. We do not penalize the coefficients on the treatment variables, which include $\{\beta_2, \beta_3, \beta_4, \beta_5, \eta_1, \eta_2, \eta_3\}$. For the other coefficients, we choose a regularization parameter using cross-validation.

Here, we assess the out-of-sample prediction accuracy of this model using cross-fitting. The sample is divided in K folds, and the prediction accuracy for farmers in fold k is assessed by estimating the model parameters on all farmers except k. This is repeated K times to estimate the predictions for farmers in every fold. The prediction accuracy and model fit of this alternative specification is compared with Equation 1. The AUC is comparable for the two specifications, 0.683. The ROC and calibration plots are provided below.

Figure A13: Out-of-Sample Predictive Performance

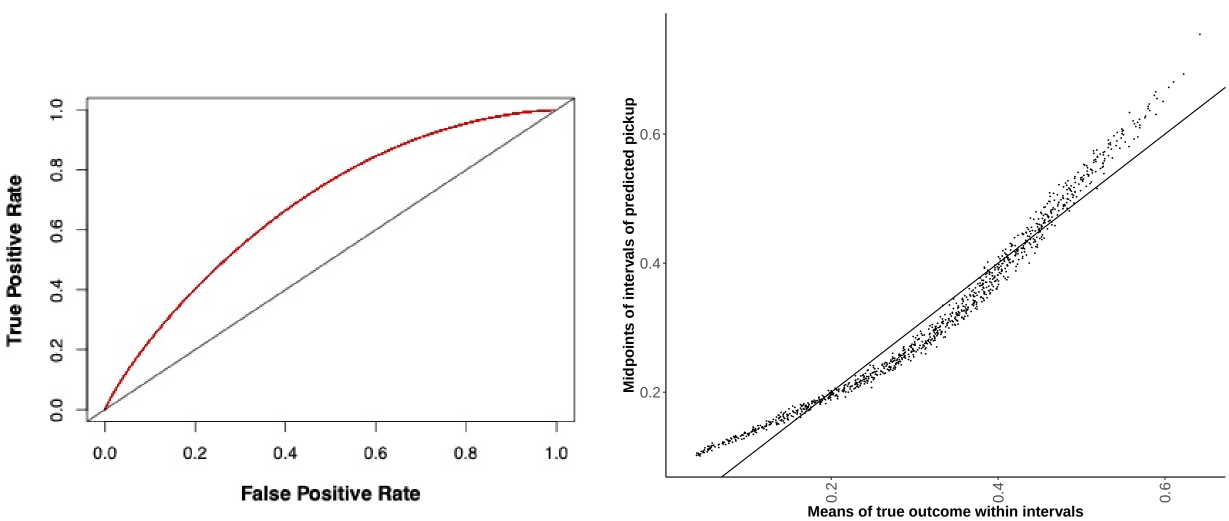

A. ROC curve B. Calibration Plot

Notes: (a) shows the ROC curve for the out-of-sample prediction. The AUC is 0.683. (b) shows the calibration plot of true and predicted pickup on the test dataset.

F.2 Hierarchical LASSO

In this section, we estimate a hierarchical LASSO model. First, Equation 5 is estimated. In this specification, the regularization parameters vary for the main effects and interaction of the covariates with the treatment dummies. Additionally, a constraint is introduced such that the penalty on the interactions is larger than the penalty on the main effects.

$$\begin{aligned}\text{logit}(p_{ij}) =& X_i\beta + \sum_j \delta_{1j}\text{Hour}_j + \sum_j \delta_{2j}\text{day}_j + \sum_j \gamma_{1j}X_i\text{hour}_j + \\ & \sum_j \gamma_{2j}X_i\text{day}_j + \sum_j \gamma_{3j}X_i\text{hour}_j\text{day}_j \end{aligned} \tag{5}$$

The objective of the hierarchical LASSO is to minimize the following-

$$\begin{aligned} -l(\theta, X, \text{hour}, \text{day}) + \lambda_1\|\beta\| + \lambda_2\|\gamma_1\| + \lambda_2\|\gamma_2\| + \lambda_3\|\gamma_3\| \\ \text{s.t.} \\ \lambda_1 \leq \lambda_2 \leq \lambda_3 \end{aligned} \tag{6}$$

We set up the optimization problem as stated in Equation 6 and solve for the regularization parameters that minimize the objective and satisfy the constraint. The *nloptr* package in R is used for optimal penalty parameters. The out-of-sample accuracy for this specification is compared with the baseline model. The AUC is 0.68, which is comparable to the AUC of the baseline specification, where we use the same penalty for all the variables except the ones on the treatment dummies.

Next, the method is modified to allow for the regularization parameters to vary by the degree of the polynomial terms and the interaction coefficients. The hour variable in the above setup has polynomial terms of degree one, two, and three, and there are interactions of the hour variable with day and other covariates. Consequently, there are five regularization parameters $\lambda_1, ..., \lambda_5$, and the objective of LASSO is to minimize the following.

$$\begin{aligned} -l(\theta, X, T) + &\lambda_1\|\beta_1\| + \lambda_2\|\gamma_1\| + \lambda_2\|\gamma_4\| + \lambda_3\|\gamma_2\| + \lambda_3\|\delta_1\| + \lambda_4\|\gamma_3\| + \lambda_4\|\delta_2\| + \lambda_5\|\delta_3\| \\ &\text{s.t.} \\ &\lambda_1 \leq \lambda_2 \leq \lambda_3 \leq \lambda_4 \leq \lambda_5 \end{aligned} \tag{7}$$

The penalty on the coefficients increases with the increasing degree of the interactions of the coefficients with treatment variables. The out-of-sample accuracy of this model is comparable to the one we estimated in Equation 1.

F.3 Matrix Completion for Historical Data

We describe the matrix completion exercise of the historical data in this section. Note that the historical data for the experimental sample is available from July 2018 to September 2021. However, we do not have the historical engagement for every farmer for each of the 91 call times, as calls were scheduled in an ad hoc manner in the past. The non-missing entries constitute about 49% of the total number of entries. Hence, a matrix completion algorithm

Figure A14: Nuclear norm penalty for the Matrix Completion on Past Engagement

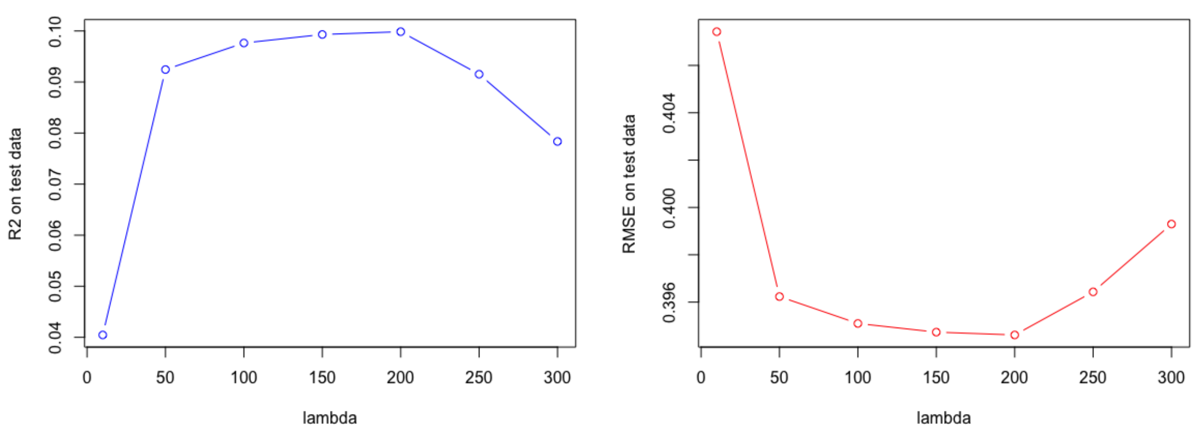

a. Out of Sample R^2 b. RMSE

Notes: (a) shows the R^2 on the test data. (b) shows the RMSE corresponding to different values of the nuclear norm penalty.

is needed to complete the past engagement matrix. The Soft-Impute algorithm described in Mazumder et al. (2010) is used for the matrix completion exercise. This algorithm uses nuclear norm regularization for matrix completion.

As described in Mazumder et al. (2010), the optimization problem for completing the historical engagement matrix Y_{history} is provided below. Note that the matrix Y_{history} has dimension $N \times J$ where each (i, j) element corresponds to the historical engagement of farmer i in call time j.

$$\min_{M} \frac{1}{2}||P_{\Omega}(Y_{\text{history}} - M)||_F^2 + \lambda||M||_*$$

, where $||M||_*$ is the sum of singular values of M and $P_{\Omega}(Y_{\text{history}})$ is the projection matrix with observed elements in Y_{history}.

The regularization parameter is chosen using the following steps. The matrix elements are split 80:20 into training and test entries. The non-missing entries in the training matrix that are part of the test IDs are made NAs for this exercise. We choose several values of λ and do the matrix completion using Soft-Impute. Both the R-MSE and out-of-sample R^2 were computed for the test entries. The λ corresponding to the lowest R^2 was chosen as the optimal λ. Panels (a) and (b) of Figure A14 display these measures for the different values of λ. The matrix completion exercise provides a 10% improvement over imputing the missing entries with the mean of the matrix. The predicted past engagement is used as a covariate for estimating and evaluating targeted policies on the uniform randomized data.