

‘Mobile’izing Agricultural Advice: Technology Adoption, Diffusion and Sustainability *

Shawn A. Cole and A. Nilesh Fernando †

April 29, 2016

Abstract

We examine the role of management in agricultural productivity, by evaluating a mobile-phone based agricultural advice service provided to farmers in India. Demand for advice is high; and advice changes practices, increasing yields in cumin (28%) and cotton (8.6%, for a sub-group receiving reminders). Information spreads, as non-treated farmers with more treated peers change practices and lose less to pest attacks. Though willingness to pay for the service is low, the value of the information externality exceeds the subsidy that would be necessary to operate the service. We estimate each dollar spent on the service yields a \$10 private return.

JEL Classification Numbers: O12, O13, Q16

Keywords: Technology Adoption, Agricultural Extension, Informational Inefficiencies

*This paper was previously circulated as ‘The Value of Advice: Evidence from the Adoption of Agricultural Practices’

†Harvard Business School (scole@hbs.edu) and Harvard University (afernando@fas.harvard.edu). This research was funded by research grants from USAID and AusAID; Cole also acknowledges support from the Division of Faculty Research and Development at Harvard Business School. Fernando acknowledges support from the HKS Sustainability Science program. In 2015, Cole co-founded a non-profit organization seeking to further test, and bring to scale, mobile-phone based agricultural advice services; this non-profit funded a post-doctoral position for Fernando. We thank the Development Support Centre (DSC) in Ahmedabad and especially Paresh Dave, Natubhai Makwana and Sachin Oza for their assistance and cooperation. Additionally, we thank Neil Patel and Awaaz.De for hosting the AO system and the Centre for Micro Finance (CMF), Chennai and Shahid Vaziralli for support in administering the evaluation. We thank Niharika Singh, Niriksha Shetty, Ishani Desai, Tanaya Devi and HK Seo for excellent research assistance. We are especially indebted to Tarun Pokiya for suggestions, his agricultural expertise and management of the AO system. This paper has benefited from comments from seminar participants at Harvard, MIT, NEUDC, and IFPRI.

1 Introduction

Agricultural productivity varies dramatically around the world. For example, India is the second largest producer of cotton in the world, after China. Yet, Indian cotton productivity ranks 78th in the world, with yields only one-third as large as those in China. While credit constraints, missing insurance markets, and poor infrastructure may account for some of this disparity, a variety of observers have pointed out the possibility that suboptimal agricultural practices may also be to blame (Jack, 2011).

This is not a novel idea. For decades, the Government of India, like most governments in the developing world, has operated a system of agricultural extension, intended to spread information on new agricultural practices and technologies, through a large work force of public extension agents. However, evidence of the efficacy of these extension services is quite limited. In India, dispersed rural populations, monitoring difficulties and a lack of accountability hamper the efficacy of traditional extension systems: fewer than 6% of the agricultural population reports having received information from these services.¹

This paper examines whether the introduction of a low-cost information and communications technology (ICT), able to deliver timely, relevant, and actionable information and advice to farmers at dramatically lower cost than any traditional service can improve agricultural management. We evaluate Aavaaj Otalo (AO), a mobile phone-based technology that allows farmers to call a hotline, ask questions and receive responses from agricultural scientists and local extension workers. Callers can also listen to answers to questions posed by other farmers.

Working with the Development Support Centre (DSC), an NGO with extensive experience in delivering agricultural extension, the research team randomly assigned toll-free access to AO to 800 households; half of this group was offered an annual traditional extension session to complement the AO service; a further 400 households served as a control group. The households were spread across 40 villages in Surendranagar district in Gujarat, India, and randomization occurred at the household level.

The AO service also included weekly push content, delivering time-sensitive information such as weather forecasts and pest planning strategies directly to farmers. This paper presents the results using three rounds of household surveys: a baseline, a midline one year later, and an endline two years after the study began. To capture information spillovers, all respondents at baseline were asked to identify individuals with whom they discussed farming:

¹This estimate is from the 59th round of the National Sample Survey (NSS) and asks farmers about their information sources for ‘modern agricultural technologies’. See Glendenning *et al.*, 2010 for a detailed discussion of this data.

we surveyed this peer group by phone in May-November of 2012.

Demand for agricultural information is substantial: more than 80% of the treatment group called into the AO line over two years. The average treatment respondent made 20 calls and used the service for more than 2.5 hours. We show that AO had a range of important, positive effects on farmer behavior. It significantly changed farmers' sources of information for sowing and input-related decisions. In particular, farmers relied less on commissions-motivated agricultural input dealers for pesticide advice and less on their prior experience for fertilizer-related decisions. Instead, farmers dramatically increase their usage of and trust in mobile phone-based information across a number of agricultural decisions.

Importantly, treated farmers were significantly more likely to adopt agricultural practices and inputs recommended by the service. These inputs choices include recommended seed varieties, fertilizers, pesticides and irrigation practices. While yields are often difficult to measure, we do find evidence of improvements here as well: treatment effects up to 8.6% higher for cotton, and 28.0% higher for cumin.

Our treatment also induces variation in information available in social networks, allowing us to estimate peer effects. We find that individuals with more treated peers (social network members) plant more cumin and report lower pest-related cotton losses. In addition, treated respondents with treated peers are more likely to adopt pest management recommendations for cotton.

Following the endline, we conduct a series of willingness to pay experiments to estimate demand for AO. Average willingness to pay for a 9-month AO subscription across multiple price elicitation methods is roughly \$2, compared to a cost of provision of for the same period of \$7. This implies a \$12 subsidy per farmer to run the service over two years. In support of such a subsidy, we estimate that each *dollar* invested in AO generates a return of more than \$10, with the return for a two-year subscription at more than \$200. Moreover, we find that a conservative estimate of the positive externality created by a treated farmer is more than 2 times the subsidy needed to cover operational costs.

First, this paper contributes to an understanding of the mechanisms underlying the dramatic variation in the productivity of firms and farms in developing countries, and the role of management practices in improving productivity. These large productivity differences have in part motivated the recent literature on non-aggregative growth (Banerjee and Duflo, 2005; Hsieh and Klenow, 2009). While a large literature focuses on the microeconomics of technology adoption (for a survey, see Foster and Rosenzweig, 2010), we instead focus on whether a consulting-like service can facilitate improved production practices. (Duflo *et al.*, 2011.) Our treatments differ from much previous work in this space in that participants receive a continuous flow of demand-oriented information, rather than a one-off provision of

supply-driven information. See [McKenzie and Woodruff, 2012](#) for a discussion of training and consulting evidence for small firms in developing countries.

More specifically, this paper advances the literature on the efficacy of agricultural extension ([Feder *et al.*, 1987](#); [Gandhi *et al.*, 2009](#); [Duflo *et al.*, 2011](#)). The existing literature finds mixed evidence of efficacy, though it is not clear whether this is due to variation in programs offered, or methodological challenges associated with evaluating programs without plausibly exogenous variation ([Birkhaeuser *et al.*, 1991](#)). This paper complements recent evidence on the historical efficacy of agricultural extension in promoting the adoption of new agricultural technologies in India ([Bardhan and Mookherjee, 2011](#)), and provides guidance as to lower-cost solutions for delivering advice. To our knowledge, our study is the first rigorous evaluation of mobile phone-based extension and, more generally, the first evaluation of a demand-driven extension service delivered by any means. [BenYishay and Mobarak \(2013\)](#) compare the impact of incentivized extension agents to non-incentivized extension agents in Malawi.

We demonstrate that informational inefficiencies are real and farmers are aware they lack information: there is considerable demand for high-quality agricultural information². Perhaps most importantly, we demonstrate that ICTs can lead to productivity gains in an industry that is both important—agriculture in the primary activity of the 47% of the world living in rural areas—and historically virtually unreliant on ICTs. Our results complement recent work that measures productivity-enhancement from ICTs in developed countries ([Draca *et al.*, 2006](#)).

We provide some evidence of the existence of a “digital divide,” as richer individuals are more likely to use the service and adopt recommended practices for cotton cultivation. This is true even though the treatment group is relatively homogeneous, and even though the technology was delivered for free, and specifically designed to be accessible to a poor, illiterate population.

Finally, we make a methodological and an empirical contribution related to measuring demand and understanding welfare among different price regimes. First, we measure willingness to pay, using two different methodologies, in a real-world setting, providing one of the first large-scale, high-stakes demonstrations that the Becker-DeGroot-Marschak (BDM) mechanism can be effective. Second, using the demand curve estimated with the BDM method, we find that a 91% subsidy induces nearly 20 times the social benefit as pricing just above marginal cost. This latter finding supports the role of subsidies in promoting the adoption of an agricultural technology in a manner similar to [Cohen *et al.* \(2010\)](#).

²Informational inefficiencies in the context of technology adoption have been defined as a situation in which farmers may not be aware of new agricultural technologies, or how they should be utilized ([Jack, 2011](#))

This paper is organized as follows. The next section provides context and the details of the AO intervention. Section 3 presents the experimental design and the empirical strategy, while Section 4 presents the results from the two years of survey data. Following this, Section 5 considers threats to the validity of the results, and Section 6 concludes.

2 Context and Intervention Description

2.1 Agricultural Extension

According to the World Bank, there are more than 1 million agricultural extension workers in developing countries, and public agencies have spent over \$10 billion dollars on public extension programs in the past five decades (Feder, 2005). The traditional extension model, “Training and Visit” extension, has been promoted by the World Bank throughout the developing world and is generally characterized by government-employed extension agents visiting farmers individually or in groups to demonstrate agricultural best practices (Anderson and Birner, 2007). Like many developing countries, India has a system of local agricultural research universities and district level extension centers, producing a wealth of specific knowledge. In 2010, the Government of India spent \$300 million on agricultural research, and a further \$60 million on public extension programs (RBI, 2010).

Yet, traditional extension faces several important challenges that limit its efficacy.

Spatial Dimension: Limited transportation infrastructure in rural areas and the high costs of delivering information in person greatly limit the reach of extension programs. The problem is particularly acute in interior villages in India, where farmers often live in houses adjacent to their plots during the agricultural cycle, creating a barrier to both the delivery and receipt of information.

Temporal Dimension: As agricultural extension is rarely provided to farmers on a recurring basis, the inability of farmers to follow-up on information delivered may limit their willingness to adopt new technologies. Infrequent and irregular meetings limit the ability to provide timely information, such as how to adapt to inclement weather or unfamiliar pest infestations.

Institutional Rigidities: In the developing world, government service providers often face institutional difficulties. The reliance on extension agents to deliver in-person information is subject to general monitoring problems in a principal-agent framework (Anderson and Feder, 2007). For example, monthly performance quotas lead agents to target the easiest-to-reach farmers, and rarely exceed targets. Political capture may also lead agents to focus outreach on groups affiliated with the local government, rather than to marginalized groups for whom

the incremental benefit may be higher. Even when an extension agent reaches farmers, the information delivered must be locally relevant, and delivered in a manner that is accessible to farmers with low levels of literacy.

The importance of these constraints is difficult to overstate (Birkhaeuser *et al.*, 1991; Saito and Weidemann, 1990.) A recent nationally representative survey shows that just 5.7% of farmers report receiving information about modern agricultural technologies from public extension agents in India (Glendenning *et al.*, 2010.) This failure is only partly attributable to the misaligned incentives of agricultural extension workers; more fundamentally, it is attributable to the high cost of reaching farmers in interior rural areas.

Finally, a potential problem is that information provision to farmers is often “top-down.” This may result in an inadequate diagnosis of the difficulties currently facing farmers, as well as information that is often too technical for semi-literate farming populations. This problem may affect adoption of new technologies as well as optimal use of current technologies.

In the absence of expert advice, farmers seek out agricultural information through word of mouth, generic broadcast programming, or agricultural input dealers, who may be poorly informed or face incentives to recommend the wrong product or excessive dosage (Anderson and Birner, 2007).³

These difficulties combine to limit the reliable flow of information from agricultural research universities to farmers, and may limit their awareness of and willingness to adopt new agricultural technologies. Overcoming these “informational inefficiencies” may therefore dramatically improve agricultural productivity and farmer welfare. The emergence of mobile phone networks and the rapid growth of mobile phone ownership across South Asia and Sub-Saharan Africa has opened up the possibility of using a completely different model in delivering agricultural extension services.

2.2 Avaaj Otalo: Mobile Phone-Based Extension

Roughly 50% of the Indian labor force, or 250 million people, are engaged in agriculture. As approximately 48% own a mobile phone (as of 2015), mobile phone-based extension could serve as many as 120 million farmers nationally⁴. Mobile phone access has fundamentally

³An audit study of 36 input dealerships in a block near our study site provides a measure of the quality of advice provided by commissions-motivated input dealers. The findings suggest that the information provided is rarely customized to specific pest management problems of the farmer, and often takes the form of ineffective pesticides that were traditionally useful, but are no longer effective against the dominant class of pests that afflict cotton cultivation.

⁴These figures are calculated using the Annual Report of the Telecom Regulatory Authority of India (India, 2015) and the World Bank Development Indicators (Group, 2012). The WDI estimates the rural population of India at 876 million while the TRAI estimates the number of subscriptions in rural India at 423 million. In addition the WDI estimate that 50% of the workforce are engaged in agriculture of a total workforce of 497 million

changed the way people communicate with each other, and has increased information flows across the country's diverse geographic areas. As coverage continues to expand in rural areas, mobile phones carry enormous promise as a means for delivering extension to the country's numerous small and marginal farmers (Aker, 2011).

Our intervention utilizes an innovative information technology service, Aavaaj Otalo (AO). AO uses an open-source platform to deliver information by phone. Information can be delivered to and shared by farmers. Farmers receive weekly push-content, which includes detailed agricultural information on weather and crop conditions that are delivered through an automated voice message.

Farmers can also call into a toll-free hotline that connects them to the AO platform and ask questions on a variety of agricultural topics of interest to them. Staff agronomists at the Development Support Centre (DSC) – our field partner – with experience in local agricultural practices receive these requests and deliver customized advice to these farmers, via recorded voice messages. Farmers may also listen and respond to the questions their peers ask on the AO platform, which is moderated by DSC. The AO interface features a touch-tone navigation system with local language prompts, developed specifically for ease of use by semi-literate farmers. The platform, which has now been deployed in a range of domains, was initially developed as part of a Berkeley-Stanford research project on human-computer interaction, in cooperation with the DSC in rural Gujarat (Patel *et al.*, 2010).

Mobile phone-based extension allows us to tackle many problems associated with traditional extension. AO has the capability to reach millions of previously excluded farmers at a virtually negligible marginal cost. Farmers in isolated villages can request and receive information from AO at any point during the agricultural season, something they are typically unable to do under traditional extension. Farmers receive calls with potentially useful agricultural information on their mobile phones, and need not leave their fields to access the information. In case a farmer misses a call, she can call back and listen to that information on the main line. AO thus largely solves the spatial problems of extension delivery discussed earlier.

A considerable innovation of AO is tackling the temporal problem of extension delivery. The agricultural cycle can be subject to unanticipated shocks such as weather irregularities and pest attacks, both of which require swift responses to minimize damage to a standing crop. Because farmers can call in and ask questions as frequently as they want, they can get updated and timely information on how to deal with these unanticipated shocks. This functionality may increase the risk-bearing capacity of farmers by empowering them with access to consistent and quality advice.

With respect to problems of an institutional nature mentioned earlier, AO facilitates

precise and low-cost monitoring. The computer platform allows easy audits of answers that staff agronomists offer, greatly limiting the agency problem. Additionally, the AO system allows for demand-driven extension, increasing the likelihood that the information is relevant and useful to farmers. Push-content is developed by polling a random set of farmers each week to elicit a representative set of concerns. In addition to this polling, the questions asked by calling into AO also provide the information provider a sense of farmers' contemporaneous concerns. This practice of demand-oriented information provision should improve both the allocation and the likelihood of utilization of the information.

However, while AO overcomes many of the challenges of traditional extension, it eliminates in-person demonstrations, which may be a particularly effective way of conveying information about agricultural practices. As discussed in the following section, our study design allows us to estimate the extent to which in-person extension serves as a complement to AO-based extension, by providing a subset of farmers with both traditional extension administered through staff at DSC and toll-free access to AO.

3 Experimental Design & Empirical Strategy

Two administrative blocks⁵, Chotila and Sayla, in the Surendranagar district of Gujarat were chosen as the site of the study, as our field partner, DSC, had done work in the area. Farmers lists, consisting of all households which expressed willingness to participate, grew cotton, and owned a mobile phone, were created in 40 villages, and served as our sampling frame.

A sample of 1,200 respondents was selected at random from this pool, with 30 households in each village participating in the study. Figure 1 summarizes the experimental design used in this study. Treatments were randomly assigned at the household-level using a scratch-card lottery. The sample was split into three equal groups. The first treatment group (hereafter, AOE) received toll-free access to AO in addition to traditional extension. The traditional extension component consisted of a single session each year lasting roughly two-and-a-half hours on DSC premises in Surendranagar. The second treatment group (hereafter, AO) received toll-free access to AO, but no offer of traditional agricultural extension, and the final set of households served as the control group. In addition, among the two treatment groups (AO and AOE), 500 were randomly selected to receive bi-weekly reminder calls (hereafter, reminder group) to use the service while the remaining 300 did not.

Figure 2 provides a timeline for the study. Baseline data was collected in June and July,

⁵A block is an administrative unit below the district level

2011, and a phone survey consisting of 798 respondents was completed in November 2011.⁶ The midline survey was completed by August 2012, and the endline survey was completed by August 2013.

To gauge balance and describe our first stage, we compute a simple difference specification of the form:

$$y_{iv} = \alpha_v + \beta_1 \text{Treat}_{iv} + \varepsilon_i \quad (1)$$

where, α_v is a village fixed effect, Treat_{iv} is an indicator variable that takes on the value 1 for an individual, i , in village v assigned to a treatment group and 0 for an individual assigned to the control group. We report robust standard errors below the coefficient estimates.

Because of random assignment, the causal effect of the intervention can be gauged by computing a standard difference-in-difference specification:

$$y_{ivt} = \alpha_v + \beta_1 \text{Treat}_{iv} + \beta_2 \text{Post}_t + \beta_3 (\text{Treat} * \text{Post})_{ivt} + \varepsilon_i \quad (2)$$

where, α_v and Treat_{iv} are as above, Post_t is an indicator variable that takes on a value of 1 if the observation was collected at the endline (or the midline) 0 otherwise, and $(\text{Treat} * \text{Post})_{ivt}$ is the interaction of the preceding two terms.

In addition, we explore heterogeneity in the treatment effect by interacting the difference-in-difference specification in Equation (2) with a dummy variable capturing the heterogeneity of interest:

$$\begin{aligned} y_{ivt} = & \alpha_v + \beta_1 \text{Treat}_{iv} + \beta_2 I(X_{iv} > \text{Median}) + \beta_3 \text{Treat}_{iv} * I(X_{iv} > \text{Median}) \\ & + \beta_4 \text{Treat}_{iv} * \text{Post}_t + \beta_5 I(X_{iv} > \text{Median}) * \text{Post}_t \\ & + \beta_6 \text{Treat}_{iv} * I(X_{iv} > \text{Median}) * \text{Post}_t + \varepsilon_{ivt} \end{aligned} \quad (3)$$

where, X_{iv} is the variable across which we explore heterogeneity in treatment effects, and $I(X_{iv} > \text{Median})$ a dummy equal to one when the observation is above the median level of X_{iv} .

While dramatically increasing statistical power, the decision to randomize at the household rather than village level raises the possibility that the control group may also have

⁶The previous version of this paper (Cole and Fernando, 2012) analyzed treatment effects using results from this phone survey.

access to information through our treatment group. This suggests that any treatment effects may in fact underestimate the value of the service.⁷ Our design asked farmers at baseline to identify peers—to measure spillovers, we subsequently collect information on 1523 peers of study respondents using a phone survey in March 2012 and November 2012, hereafter the ‘peer survey’.⁸ This data allows us to estimate whether the treatment also influences the outcomes of individuals in our study respondents’ social networks. We estimate the extent of such peer effects or information spillovers with the following specification:

$$y_{iv} = \alpha_v + \beta \left(\frac{\# \text{References in Treatment}}{\# \text{References}} \right)_{iv} + \sum_{i=2}^7 I(\# \text{References} = i)_{iv} + \varepsilon_{iv} \quad (4)$$

where, α_v is as above, $\sum_{i=2}^7 I(\# \text{References} = i)_{iv}$ is a fixed effect for the number of peers who cite a respondent as a top agricultural contact and $\left(\frac{\# \text{References in Treatment}}{\# \text{References}} \right)_{iv}$ is the fraction of these respondents who are assigned to treatment.

We did not prepare a pre-analysis plan prior to undertaking the study. This was in part due to the dynamic nature of the treatment: the service responded to farmer questions, and ex-ante, it was not always clear which subjects farmers would inquire about. We address concerns about multiple inference in two ways. First, we use the content generated by farmers, and by our agronomist, as a broad guide for conducting empirical analysis.⁹ Second, we aggregate agricultural practices into indices, following, for example, [Kling *et al.* \(2007\)](#).

To construct indices, we do the following. The agronomist in our study, Tarun Pokiya, characterized the full set of agricultural practices that were recommended by the service. We then aggregate all variables corresponding to recommended practices by calculating a z-score for each component and then take the average z-score across components. Each component z-score is computed relative to the control group mean and standard deviation at baseline. We have compared this to the method that uses ‘seemingly unrelated regression’ which gives slightly different standard errors and identical point estimates but is virtually indistinguishable from this method as suggested by [Kling *et al.* \(2007\)](#).

⁷In order to control for spillovers, we estimated the main difference-in-differences specification with controls for the fraction of a respondent’s social network that was also a part of the study and the fraction that was assigned to treatment. These controls leave the main estimates largely unchanged.

⁸At baseline, we asked all respondents to list the three contacts with whom they most frequently discussed agricultural information and collected their phone numbers. The ‘peer survey’ collected information from all these contacts. Note, some of these 1523 peers may themselves be study respondents. The analysis largely focuses on 1114 non-study peers

⁹See Appendix Table A1 for details of questions asked by farmers on the AO service and push content provided.

3.1 Summary Statistics and Balance

In this section we assess balance between the ‘combined treatment’ group (AO + AOE) and the control group and the subset of the treatment group that receives reminder calls, referred to as the ‘reminder’ group, and the control group. We do not find many differences in the separate treatment effects of the AO and AOE groups and the interaction of these treatments with reminder calls; the exposition of our paper therefore focuses on the ‘combined treatment’ group, as well as the ‘reminder’ group.¹⁰

Table 1 contains summary statistics for age, education, income and cultivation patterns for respondents in the study, using data from a baseline paper survey conducted in July and August of 2011. Column (1) reports the mean and standard deviation for the control group, column (2) tests the initial randomization balance between the combined treatment group and the control group. Finally, column (3) tests the balance between the reminder group and the control group.

We see that respondents are on average 46 years old and have approximately 4 years of education. Columns (2) and (3) show that the randomization was largely successful for both treatment groups across demographic characteristics (Panel B) and indices capturing information sources, crop-specific and general input use (Panel C). However, an imbalance exists in the area of cotton planted between the treatment groups and the control group in 2010 but not in 2011 (both periods are prior to treatment).¹¹ The combined treatment group is also more likely to grow wheat, but this crop is mostly grown for home consumption in this context.

As cotton is the most important crop in our sample, we take a conservative approach to the possibility that baseline cotton levels affect subsequent outcomes and include as controls the area of cotton cultivated in 2010 and its interaction with the ‘Post’ term in both the difference-in-difference specification (equation (2)), the heterogeneous effects specification (equation (3)) and peer effects specification (equation (4)).¹²

¹⁰Appendix A2 tests the balance for the AO and AOE group, while Appendix A3 reports treatment effects separately for the AO and the AOE group.

¹¹Note, the 2011 figures for wheat and cumin are not reported as they are grown during the Rabi season after the treatment was administered.

¹²Appendix Table A4 provides a more systematic treatment of balance in our sample. We look for significant differences in baseline characteristics between the combined treatment group and control, and the reminder group and control respondents. Among the differences computed using the latter specification (examining all 2,295 baseline variables) we find that 0.7% are significantly different from zero at the 1% level, 4.4% are different at the 5% level of significance and 9.5% at the 10% level. These results confirm that the randomization was successful, and that the cotton imbalance is a result of chance rather than any systematic mistake in the randomization mechanism.

4 Experimental Results

Cole and Fernando, 2012 describe initial differences measured seven months after the implementation of AO. After seven months, take-up among the treated group was high, and we measured several important changes in agricultural behavior: farmers changed their information-gathering activity, relying less on peers and more on mobile phone-based advice; treatment farmers were more likely to adopt more effective pesticides, and reduce expenditure on hazardous, ineffective pesticides; and treated farmers were more likely to grow cumin. A short-coming of the early evidence was that it was based on interviews, conducted by telephone, of only a sub-sample of study participants. In the sections below, we describe results after treatment households had been offered the service for two full years primarily using the difference-in-differences specification in Equation (2).

4.1 First Stage: Take-Up and Usage of AO

Table 2 reports information on take-up and usage (first stage). While control respondents were not barred from AO usage, only four control respondents called into the AO line by the midline and a further 25 had called in after two years. As a result, virtually all AO usage is accounted for by respondents in the treatment group. As of August 2013, two years after commencement of the service, 88% of the treatment group had called into the AO line, making an average of 22 calls. This represents a substantial increase from the midline, where 65% of the combined treatment group called in, making an average of 9 calls. The mean usage for treatment respondents is over 2.5 hours, as compared to 1.3 hours at midline. On average, treatment respondents have listened to 53% of total push call content (54% of total push call content was the average at the midline). By the endline, the average number of questions asked by the treatment group is 1.7, with 9% of the treatment group responding to a question. Further, columns 4 (midline) and 8 (endline) show that the reminder group had used the service almost an hour more on average, but were not statistically more likely to call into the line.

Taken together the results represent substantial induced usage for treatment farmers, although one-fifth of the treatment group did not use the service. Additionally, these average effects also mask important temporal patterns shown in Figure 3 which reports average AO use by month. We see that there was substantial usage across treatment arms during the first six months after the intervention was administered. Following this period, usage has been trending down, but with important spikes during sowing times and harvest time. This figure is suggestive of AO users acquiring a stock of knowledge and supplementing thereafter with dynamic information needs throughout the season.

Appendix Table A1 provides a categorization of the questions asked by treatment respondents during the two years of service. (The categories are not mutually exclusive.) Unsurprisingly, columns 3 and 4 show that most questions (50%) relate to cotton, and a majority (54%) focus on pest management and these numbers are relatively stable across both years. Table A1 also reports information on the content of push calls (columns 5-8), which tended to provide more information on cumin and wheat cultivation than incoming questions and were the primary source for weather information.

4.2 Impact on Sources of Information for Agricultural Decisions

Panel A of Table 3 examines the use of mobile phone based information in agricultural decision-making, and measured trust (on a scale of 1-10) of information provided by mobile phones. By the endline, treatment farmers are 70 percent more likely to report using mobile phone-based information to make agricultural decisions. The treatment effect on reported level of trust in mobile phone-based information is also dramatically higher: approximately 6.27 points greater on a 10-point scale. An index aggregating the importance of mobile phone based information (analysis of the topics comprising this index follows immediately below) for all subject areas is 1.26 standard deviations higher in the treatment group.

We asked farmers for their most important source of information for a series of agricultural decisions. The survey responses are recorded as free text, without prompting, and coded into categories by our data entry teams. We present results across a variety of subject areas. Panel B of Table 3 shows that the treatment group consistently reports using mobile phone-based information across a series of agricultural decisions. By the endline, large effect sizes can be seen in the case of pest management (24.3%) and smaller effects in the case of fertilizer decisions (10%) and crop planning (5.6%).

Other than input-related decisions, mobile phone information is used increasingly by the treatment group for other topics such as weather (36.8%). Importantly, we do not find any effect of our treatment on the use of mobile phones for price information. The AO service never provided price information. This helps address the concern that social desirability bias may be contributing to our results. Additionally, across virtually all agricultural decisions, we do not observe statistically or economically significant differences between the combined treatment group and the reminder group.

Appendix A5 provides more disaggregated effects on sources of information. As suggested by the index of information sources, we observe across the board increases in the use of mobile phone-based information. The treatment group reports using information from input dealers less often in making pesticide decisions (-7.2% at midline), although, interestingly, they report consulting input dealers *more* often in the case of cotton fertilizer use (5%)

and cumin planting (3.7%) at the endline. There are also reported reductions in the use of information from ‘other farmers’ and ‘past experience’. The reduction in reliance on past experience for cumin fertilizers is significant at the midline.

Taken together, these results suggest that AO has been successful in establishing itself as a source of information for treatment respondents in making a variety of important agricultural decisions. These results also suggest that demand exists for agricultural information in rural Gujarat and that this information is not currently being provided via mobile phone. In the next sections we look at whether the provision of information through AO affected input use and agricultural productivity more broadly.¹³

4.3 Overall Impact on Input Adoption

A number of input choices influence agricultural productivity. Cotton is the main cash crop grown in our sample – grown by 98.4% of the sample at baseline – and chemical inputs such as pesticides and fertilizers greatly affect cotton yields.¹⁴ In addition, Bt cotton is the dominant variety of cotton grown in this context – although there are literally hundreds of sub-varieties and brands which pose other difficulties – and yields are particularly sensitive to regular irrigation.

Panel A of Table 4 shows that total input expenditure is not significantly different between the combined treatment group and the control group at either the midline or the endline. However, we observe that expenditure on irrigation is twice as high for the combined treatment group and the reminder group at the midline (significant at the 1% level). Similarly, by the endline, irrigation is 60% higher in the combined treatment group (t-statistic = 1.64) and 80% higher in the reminder group (significant at the 5% level).¹⁵ Irrigation was not a leading topic for our service (we received 21 questions about it, and covered it in five push calls), it is possible that farmers felt more confident spending resources on irrigation if they believed other risks would be easier to address because of the information service.

Panel B of Table 4 shows that the treatment group consistently adopted more cotton-related inputs and practices suggested by the service (0.05-0.07 standard deviation units). These input decisions include recommended seed varieties, pesticides, fertilizers and irriga-

¹³Appendix A2 provides even more detail on changes in sources of information. Across a number of agricultural decisions, farmers tend to rely heavily on other farmers, with input shops being particularly important for pesticide decisions. Notably unimportant are government extension services, virtually unmentioned by farmers as a source of information.

¹⁴In 2006-2007, 87% of all land under cotton in India was treated with pesticide. In contrast, this figure is just 51% for paddy and 12% for wheat. Calculations by author (Agricultural Census of India, 2006).

¹⁵Panel B of Appendix Table A6 reports a detailed breakdown of changes in input costs. In addition to changes in irrigation costs, we observe changes in expenditure on seeds, but these changes are not significant at traditional levels (t-statistic = 1.4).

tion practices. While the treatment effects on the overall wheat and cumin indices are not significantly different from zero, the point estimates are qualitatively consistent.¹⁶

4.4 Impact on Seed Selection

The presence of a wide variety of cotton seeds, some counterfeit, makes seed selection a particularly important decision. In Uganda, [Bold *et al.* \(2015\)](#) demonstrate that low quality inputs dramatically depress returns to hybrid seeds. In Panel C of Table 4, we observe that the index of cotton seed-related decisions is consistently higher (0.046-0.051 standard deviation units) in the combined treatment group and the reminder group at midline and is statistically significant at the 10% level.

An inventory analysis in Sayla and Chotila following conclusion of the study verified that many of the products we recommended were indeed commonly stocked by local input dealers. For example, we find treatment farmers purchased 0.08 kg more of Ganga Kaveri, a brand we recommended, relative to control groups.

4.5 Pest Management Practices

In Panel D of Table 4, we examine the treatment effect on pest management practices. The index which includes all pest management practices is 0.08 standard deviation units higher in the reminder group at the endline. This effect was not significant for the combined treatment group, but all estimated coefficients move in the same direction.

Examining the sub-components of the index (see Appendix Table A7), there are no statistically significant results for pesticide purchase and usage for the treatment group, once again in contrast to the simple difference estimates 7 months after the intervention has been administered in [Cole and Fernando \(2012\)](#). This early version of the paper reported a simple difference in the use of imidachlorprid but subsequent difference-in-difference analysis in this paper revealed this effect was driven by baseline imbalance.¹⁷

We do observe a 2.4% increase (2.1% in the midline) in the fraction of treatment respondents using trichoderma, a biological method of pest control, relative to the control group. The AO service provided extensive information in both Kharif and Rabi on the use of Trichoderma, as a means of preventing wilt disease in cotton and cumin.

¹⁶The standard errors also suggest that the experiment may be underpowered to detect effects for cumin (grown by just 34% of the sample), while wheat cultivation involves substantially fewer chemical inputs and is primarily produced for home consumption.

¹⁷While total money spent on acetamaprid increases, this number is only significant for the AOE group (an increase of Rs. 80, not reported). Similarly, while total spent on monocrotophos decreases, the only statistically significant result is among the AOE treatment group (a decrease of Rs. 60, not reported).

4.6 Fertilizers

In Panel E of Table 4, we examine fertilizer practices. The index of cotton fertilizer practices is 0.07 standard deviation units higher among the combined treatment group in the endline, as compared to 0.10 standard deviation units in the midline. The index is similarly higher for the reminder group at both the midline and the endline, but the point estimates are not statistically significant.

The disaggregated results (reported in Appendix Table A7) indicate farmers are purchasing more of the fertilizer our service recommended. In particular, we observe a 5.7% (4.5%) increase in purchases of ammonium sulfate at the midline (endline), and an increase of 5.5% in NPK Grade 1 fertilizer at endline. To put these effect sizes into perspective, [Dufflo *et al.* \(2011\)](#) find an increase of 16-20% in fertilizer adoption in Kenya using free delivery of planting and top-dressing fertilizer, while [BenYishay and Mobarak \(2013\)](#) find increases of 2.2-5.5% across treatments in pit planting and 0-19% across treatments for composting in a study using in-person physical extension.

4.7 Sowing and Productivity

In Table 5 we examine sowing choices and agricultural productivity. We do not observe any effect of the treatments on the frequency of cultivation or area planted of cotton, cumin or wheat.

Panel B shows that cotton yields are consistently higher for the treatment group and the reminder group at the midline and the endline. However, this effect is only significant for the reminder group at the midline (increase of 60 kg per acre, or 8.6% higher than the control mean). Additionally, we see that yield for cumin is about 48 kilograms per acre higher at endline (28% higher than the control mean) among the treatment group and 54 kilograms per acre higher for the reminder group (31.4% higher than the control mean) and statistically significant at the 5%-level. These results are robust to winsorizing ($p=0.25$).¹⁸

As in the case of agricultural yields, the detection of treatment effects on profits is greatly complicated by measurement error. At the endline, both the treatment group and the reminder group have profits that are more than \$200 higher than the control group (16% higher), although both these effects are imprecisely estimated. In addition, we see an 8% increase in input expenditure by the endline for the combined treatment group (26% higher for the reminder group), but this effect is also imprecisely estimated. Measuring in levels rather than logs, we find that input expenditure is higher for the reminder group at endline

¹⁸Note, all percentage increases of estimates of yield are computed by dividing the coefficient on $treat*post$ by the control mean at baseline.

by roughly \$50, significant at the 10% level (not reported).

4.8 Impact on Agricultural Knowledge

Having established that AO affects behavior, we now turn to the mechanisms by which AO works: does it serve as an education tool, creating durable improvements in knowledge, or does it function as an advisory service, in which farmers follow instructions, without necessarily comprehending why a particular course of action is the right one? In Table 6, we examine whether AO improves farmers’ ability to answer basic agricultural questions. The questions we ask test the respondents on a wide range of topics, which are generally invariant to their personal circumstances.¹⁹

Baseline agricultural knowledge is low, with farmers in the control group only being able to answer 32% of questions correctly. There are no imbalances between treatment and control for the total at the baseline. Given that these are very basic questions about agriculture, this suggests that there is a substantial lack of information on even basic topics concerning crop cultivation.

As reported in Table 6, we do not observe differences between the treatment and control groups in agricultural knowledge in the midline or in the endline survey. In part, the types of knowledge that respondents gain reflect their actual demand for information. The majority of questions asked on the AO platform relate to pesticides.

4.9 Heterogeneous Treatment Effects

While the importance of technological progress to growth is beyond doubt, there are growing concerns about the possibility of a “digital divide,” in which the poorest or least educated are less able to take full advantage of the promise of new technologies. We test this hypothesis by comparing AO usage and knowledge gain by education level. We focus on respondent education for at least two reasons: first, while the service is designed to be accessible to illiterate users, it may be easier to use or navigate for a literate population, who can take advantage of instructional material. Second, educated individuals may be in a better position to learn. (We also examined landholdings as a source of heterogeneity in treatment effects, and found virtually no difference between above- and below-median landholders.) The median farmer in our survey reports 4 years of education.

Are AO Usage and Education Complements?

In Table 7, we regress measures of AO usage on a treatment dummy, a dummy for having

¹⁹The full text of the questions is available in Appendix A8.

more than the median number of years of formal education (4 years), a time-trend dummy and the corresponding interaction terms as in equation (3).

Columns (2) and (3) suggest that there may be some complementarities between AO use and education: more educated farmers make more use of the service on average, but these differences are not statistically significant. We do not find an effect on the extensive margin; that is, more educated individuals are no more likely to call into the AO line. This table makes use of administrative data for all 1,200 respondents as their calls (and the absence of calls from control) are logged on to the server. We do not observe heterogeneous effects of AO across education for input adoption or agricultural knowledge.

Income and AO

Treatment respondents with above-median incomes are no more likely to call into the AO line, but their total usage is approximately 40 minutes higher relative to respondents with below-median incomes at midline (55 minutes in the endline). Farmers with higher incomes also show differential effects in the cotton practices index (about 0.1 standard deviation units higher, 0.07 at midline but not significant).

4.10 Spillover Effects

Given randomization at the household level, it is possible that access to AO indirectly influenced the outcomes of study respondents and those not a part of the study in the networks of study respondents through information spillovers.

In a separate paper, we document in detail how patterns of social interactions and information exchange are influenced by the AO treatment (Fernando, 2016). We find that the technology results in treatment respondents being 7.2% more likely to share information with their peers and 7% more likely to recommend an input after production outcomes have been observed (both significant at 5% level). In addition, they are 46.8% more likely to report ‘mobile phone-based information’ as the source of this information, suggesting that the treatment both influenced the frequency and content of information sharing (significant at the 1% level).

Table 8 estimates spillover effects for both study respondents and a group of ‘non-study’ respondents who were surveyed in the ‘peer survey’. As the previous results suggest, treatment respondents may have discussed advice they received or even asked questions on behalf of their peers. Alternatively, peers may follow suit after directly observing changes in their neighbors’ agricultural practices.

4.10.1 Among Study Respondents

In Columns (1)-(7) of Table 8 we use the same specification as in the section on heterogeneous effects (Equation (3)), except the heterogeneity explored is the fraction of one’s peer group exposed to the treatment as in Equation (4) but in a difference-in-difference framework.²⁰

Column (1) contains the mean and standard deviation for the control group at baseline, columns (2)-(4) refer to study respondents at midline, while columns (5)-(7) refer to them at endline, from a separate regression.²¹

Columns (2) and (5) report the coefficient on the interaction between an indicator for treatment and the post variable. Importantly, these estimates do not differ substantially from those estimated in previous tables, suggesting that controlling for spillover effects does not influence headline results on AO usage and reported sources of information.²² Columns (3) and (6) report the coefficient on the interaction between post and the fraction of one’s peer group that is treated, while columns (4) and (7) show the interaction of the former with an indicator for belonging to the treatment group. These coefficients are for the most part not statistically significant, providing limited evidence for spillover effects within the study. An important exception is the cotton pest management index. Here we see that treated respondents with a higher fraction of treated peers in their network are more likely to adopt pesticide recommendations at midline and at endline, although the endline difference is not statistically significant. In addition, we see some evidence that among control respondents, having more treated peers made it less likely that they would adopt recommendations for cotton pest management at midline (significant) and endline (not significant), an effect that may suggest limitations on input inventories within villages. In both cases, the net treatment effect accounting for spillovers is significantly different from zero.²³ These effects are suggestive of complementarities between treated peers in the adoption of cotton pest management advice.

4.10.2 Non-study Respondents

Columns (8) and (9) refer to non-study respondents and report simple differences using data collected from the peer survey. The specification estimated here is Equation (4), with controls

²⁰In each case, we control for the number of peers in one’s reference group, its interaction with a time-trend and with a treatment indicator, and baseline cotton and its interaction with a time-trend.

²¹Appendix A9 assesses whether the fraction of treated peers in a social network is independent of other observable characteristics. The only characteristic that shows an imbalance is cotton acreage. We control for baseline cotton acreage in all peer regressions.

²²Appendix A11 shows that this holds true for a broader set of outcomes.

²³We can reject the null hypothesis for the joint test of significance for treat*post coefficient and the treat*post*treat*frac at both midline (p-value: 0.078) and endline (p-value : 0.081) for the cotton pest management effect.

for the number of peers in one’s reference group and baseline cotton.

Here, we do find significant (economically and statistically) impacts of informational spill-overs. Non-study respondents with more treated peers are more likely to grow cumin (6%) and plant a large amount of it (.26 acres more). Those with more treated peers in their networks also report 4% less cotton crop loss as a result of pest attacks, suggesting that pest management practices provided by the AO service may have been shared.

4.11 Willingness to Pay

Most goods and services are evaluated on a market basis, rather than through RCTs. The financial sustainability of a subscription-based service would depend critically on users willingness to pay—though we point out that because information is shared among farmers, willingness to pay may well be less than the social value of the service.

After the conclusion of the study, we conducted a series of exercises to assess willingness to pay for the AO service among the original 1,200 study respondents, as well as an additional 457 non-study respondents. The first method used a traditional offer price (‘Take it or Leave it’ (TIOLI)), for a nine-month subscription to AO. We randomly varied the offer price at the household level to estimate a demand curve²⁴ The second method used the Becker-DeGroot-Marschak (BDM) method as an incentive compatible price elicitation mechanism. In this method, the respondent first indicates their willingness to purchase at a series of price points. They then record a specific bid, after which the respondent is shown a randomly generated offer price.²⁵ To our knowledge, this was one of the first successful implementations of the BDM mechanism “in the field” for a substantial product. If the respondent’s bid is greater than the offer price they can buy it at the offer price and if not they cannot purchase the product. The TIOLI method was randomized to a quarter of the sample, while the BDM method was randomized to the remaining three-quarters.

The two methods of eliciting willingness to pay deliver similar results. Of the 390 respondents that were offered AO through the TIOLI method, 150 respondents (38.4%) bought a subscription at an average price of Rs. 107 (\$1.78). Similarly, of the 1043 respondents that were offered AO through the BDM method, 370 (33%) purchased a subscription at an average price of Rs. 108 (\$1.8).²⁶

Table 9 investigates correlates of the decision to purchase AO. Surprisingly, we do not find that treatment status is an important predictor of purchasing AO. Rather, we find that education positively predicts the decision to purchase AO, while the offer price does the

²⁴The prices offered were Rs. 40 (\$0.67), Rs. 90 (\$1.5), Rs. 140 (\$ 2.3), Rs. 190 (\$3.2) and Rs. 240 (\$4).

²⁵The respondent is asked to indicate their willingness to purchase the policy for Rs. 40 (\$0.67), Rs. 90 (\$1.5), Rs. 140 (\$ 2.3), Rs. 190 (\$3.2), Rs. 240 (\$4), Rs. 290 (\$4.8), Rs. 390 (\$6.5), Rs.490 (\$8.1)

²⁶Cole *et al.*, 2016 describe and evaluate these approaches in detail.

opposite.

Figure 4 shows the elicited demand curves for AO for both methods. The methods yield comparable estimates of willingness to pay, which we estimate at Rs. 108 (\$1.78) for a nine month subscription (see Panel A of Table 9) . AO costs little, requiring just \$0.83 to service one farmer per month, inclusive of airtime costs, staff time and technology fees. In contrast, a single round of traditional extension (educational demonstration by a government extension worker to a gathering of farmers) costs \$8.5 per farmer (based on extension provided to the AOE group).

In our study, airtime was provided freely for farmers to encourage take-up (costing approximately \$ 0.31). If farmers paid airtime, the per-farmer operating cost of the AO service could be as low as \$0.52 per month. However, even at this rate AO would require a subsidy of roughly \$0.35 per month per farmer given the elicited willingness to pay. It is important to note that the per-farmer cost of providing AO is likely to drop considerably as the service scales up, as labor costs need not scale linearly if pre-recorded answers can be directed to commonly asked questions, and if information can be transmitted using cheaper data plans, rather than voice calls.

4.12 Cost-Benefit Analysis

To compute the return to investing in an AO subscription we weigh measured increases in yield against increases inputs costs. A 8.6% increase in cotton yields for the treatment group with frequent reminder calls implies an average revenue increase of nearly \$200 while a 28.0% increase in cumin yields implies an average return of \$65.²⁷ This \$265 average increase in revenue must be weighed against an increase in input costs of \$50.²⁸ This implies a profit of \$215 on the basis of a \$20, 2-year subscription to AO. This implies a return of more than \$10 for each dollar invested in AO, net of cost.

4.12.1 Accounting for Externalities

While we estimate the average private return to a respondent at \$215 per farmer, we can also compute the social return by estimating externalities caused by the service. In particular, we find that exposing a non-study peer to a treated respondent reduces crop loss due to pest

²⁷These calculations are based on average values of crop acreage and crop selling prices for the entire sample. On average, respondents grew 4.4 acres of cotton (0.55 acres of cumin) and sold cotton at a price of \$0.74 per kg (\$2.18 per kg for cumin) at the time of the endline survey. We observe an increase of 60 kg per acre in cotton yields and 54 kg per acre for cumin yields for the treatment group.

²⁸Input costs include the costs of seeds, irrigation, fertilizers, pesticides, hired and household labor. Household labor is priced at the mean of the hired wage. This effect is precisely estimated for the reminder group at the endline. The values at midline and for the combined treatment group imply a smaller increase in input costs but are not precisely estimated.

attacks by 4%. Assuming this effect is linear, this suggests that a representative treatment respondent creates a positive social externality amounting to \$16²⁹. As such, the social return, net of cost is \$231, or a return of \$11.55 for each dollar invested. While this estimate does not take into account additional costs incurred by peers in reducing crop loss, we view this as a conservative estimate given other potential benefits such as increased cumin cultivation.³⁰ Even under these assumptions, it is worth noting that this externality nearly 3 times the per-farmer subsidy cost (\$5.69 per farmer for a 9-month subscription) needed to keep the service operational given elicited average WTP.

4.12.2 Pricing and Welfare Analysis

Given the demand curve we estimate, the profit-maximizing price for a private firm would be Rs. 490 (the highest price we tested), which would yield the firm a profit of \$8,930 across 1,114 individuals, and create a social benefit of \$10,154 (\$231 per farmer * 50 farmers purchasing the service)³¹. In contrast, were the service offered at a 91% subsidy, with a subscription cost of Rs. 40, the net benefit would be \$193,589. Because the benefit is so much greater than the willingness to pay, a heavily subsidized service would generate 20 times more social benefit than pricing above marginal cost. While these calculations do not take into account distortionary effects that a tax needed to raise such a subsidy would generate, these are likely to be second order relative to the benefits generated and suggest a strong case for subsidizing the service in the interest of maximizing social welfare.

²⁹In our study, we treated 800 farmers across 40 villages with a total population of approximately 60,000. These study respondents had 1,114 unique peers who were not in the study sample and we estimate the savings per peer as 3.9% of cotton yield, or approximately \$20 per peer, assuming baseline cotton acreage of 694 kg and a price of 0.74 kg per acre. However, each peer only had 0.58 of their peers in the treatment group, so we compute the total externality as \$20* 1114 * 0.58, or the per-treated-farmer externality as \$16. Had we treated a far higher share of the population, this number may well have been lower.

³⁰For example, the treatment also increases cumin acreage by 0.225 acres. Assuming yield is linear in acreage, this amounts to 70.6 additional kg of cumin or \$153. However, cumin is grown by 43% of the sample (in contrast to 99% for cotton) by endline, and pricing these benefits requires explicit assumptions on adoption and the production function.

³¹This calculation takes the Rs. 490 price point under the BDM game, at which just 4.5% of 1,114 respondents administered the BDM game are willing to buy a 9-month subscription to AO. This calculation further assumes that the per-farmer private and social benefit of the service is \$231. The net social benefit is the difference between the surplus generated for the farmers in increased returns, both private and social, and the costs incurred by the firm in providing the service. Note, pricing just below marginal cost (Rs. 390) would result in \$15,556 of net social benefit

5 Threats to Validity

5.1 Attrition

In the endline survey, we had 120 attritees, of which 39 were control farmers, 43 from the AOE group, and 38 from the AO group. In comparison, we had 77 attritees in the midline, of which 23 were control farmers, 22 were from the AOE group and 32 were from the AO group. We do not observe any significant differences between the treatment and control group for the attritees, as measured by baseline characteristics. These results are reported in Appendix Table A10.

5.2 Experimenter Demand Effects

A second obvious concern is that respondents in the treatment group may offer answers that they believe the research team seeks, perhaps in the hopes of prolonging the research project, or due to a sense of reciprocity. While it is difficult to rule this out entirely, the fact that we find no effect on sources of price information in Table 3 – which the AO service does not provide – in spite of finding large differences for sources of other information provides some comfort. We also note that we can observe some outcomes perfectly: the AO platform records precisely how many times respondents call in. Respondents provide remarkably unbiased answers to the question “did you call into the AO line with a question,” with 55.5% self-reported call-in rate vs. a 53.5% call-in rate using administrative data (results not reported in tables).

6 Conclusion

This paper presents the results from a randomized experiment studying the impact of providing toll-free access to AO, a mobile phone-based technology that allows farmers to receive timely agricultural information from expert agronomists and their peers.

Firstly, we show that the intervention was successful in generating a substantial amount of AO usage, with roughly 60% of the treatment group calling into listen to content or ask a question within 7 months of beginning the intervention, and 80% using it after two years. We then showed that AO had a large impact on reported sources of information used in agricultural decisions, reducing the reliance of treatment respondents on input dealers and past experience for advice.

Having established AO as a reliable source of information, we then show that advice provided through AO resulted in farmers changing a wide variety of input decisions that

ultimately lead to increases in crop yields. In addition, we find evidence that treated respondents had a limited influence on the information sources and cropping decisions of peers not in the study. Richer respondents are more likely to use AO and adopt inputs, suggesting that richer farmers may be differentially well-positioned to take advantage of technological change.

We estimate that a \$1 investment in AO generates a return of more than \$10. Elicited willingness to pay for a \$7.5 subscription is only \$1.7, but implied subsidy is more than justified by the returns generated by AO. A two-year subscription generates a profit of more than \$200 on average, while inducing a positive social externality of \$16, or nearly 3 times the subsidy required to operate the service given elicited WTP. In addition, while the cost of this intervention is quite low (we estimate a monthly cost of approximately USD \$0.83 per farmer, including all airtime costs, staff time, and technology fees) if the project were implemented at scale, the costs may drop dramatically, as pre-recorded answers to specific questions dramatically reduce the amount of time the agronomists must spend on each question. In contrast, the “all-in” costs for physical extension were about \$8.50 per farmer. In addition to this high cost, we do not find any evidence to suggest that outcomes between respondents provided with AO and physical extension and those only provided with AO were different.

These results represent the beginning of a research agenda seeking to understand the importance of information and management in small farmer agriculture. Many important questions remain unanswered. Going forward, the individual nature of delivery and information access (each farmer can potentially receive a different push call message, and each can choose which other reported experiences to listen to) will allow us to test the importance of top-down vs. bottom-up information.

One of the features of the current intervention is that the NGO providing the service, DSC, has established trust by providing services to farmers for many years. While certain aspects of observed input adoption like pesticide use allow for sequential learning, for large investments where the downside risk could be potentially devastating, as in the case of cumin sowing, trust would appear to be a lot more important. AO comes across as a service without a vested interest (impartial) in addition to being experts, which may well serve to both encourage farmers to switch away from other sources and act on AO information. We hope to experimentally vary the source of information (if only to present it as a peer instead of an expert) in order to understand the importance of this aspect for technology adoption.

To understand the exact mechanism through which AO affects behavior, it is also important to understand whether the treatment effect is working through acquired knowledge or “merely” persuasion. One definition of cognitive persuasion that has been adopted in the

literature is that it consists of “tapping into already prevailing mental models and beliefs” through associations rather than teaching or inculcating the subject with new information. From qualitative work we have conducted, many farmers claim to distrust input dealerships but still adopt their advice for lack of a better source. While this is not something that is emphasized in the AO service itself, the presentation of information that seems to conflict with the advice given by input merchants may well serve to reinforce this distrust. We hope to be able to test these hypotheses using pre- and post- subjective evaluations of the trustworthiness of information sources. However, a more elaborate treatment may be necessary to clearly distinguish between the two models of how information affects behavior.

Finally, we stress the practical importance of this technology. Climate change and the mono-cropping of new varieties of cotton may significantly alter both the types and frequency of pests, and the effectiveness of pesticides in the near future. Farmers in isolated rural areas have little recourse to scientific information that might allow them to adapt to these contingencies. We believe mobile phone-based agricultural extension presents a cost-effective and salient conduit through which to relay such information.

References

- AKER, J. (2011). Dial "a" for agriculture: Using ict's for agricultural extension in development countries. *Agricultural Economics*.
- ANDERSON, J. and BIRNER, R. (2007). How to make agricultural extension demand-driven? : The case of india's agricultural extension policy. *IFPRI Discussion Paper 00729*.
- ANDERSON, K. and FEDER, G. (2007). Chapter 44: Agricultural extension. *Handbook of Agricultural Economics*.
- BANERJEE, A. V. and DUFLO, E. (2005). Growth theory through the lens of development economics. *Handbook of economic growth*, **1**, 473–552.
- BARDHAN, P. and MOOKHERJEE, D. (2011). Subsidized farm input programs and agricultural performance: A farm-level analysis of west bengal's green revolution, 1982-1995. *American Economic Journal: Applied Economics*, **3** (4).
- BENYISHAY, A. and MOBARAK, A. M. (2013). Communicating with farmers through social networks. *Yale University Economic Growth Center Discussion Paper*, (1030).
- BIRKHAUSER, D., EVENSON, R. and FEDER, G. (1991). The economic impact of agricultural extension: A review. *Economic Development and Cultural Change*.
- BOLD, T., KAIZZI, K. C., SVENSSON, J., YANAGIZAWA-DROTT, D. *et al.* (2015). *Low Quality, Low Returns, Low Adoption: Evidence from the Market for Fertilizer and Hybrid Seed in Uganda*. Tech. rep.
- COHEN, J., DUPAS, P. *et al.* (2010). Free distribution or cost-sharing? evidence from a randomized malaria prevention experiment. *Quarterly Journal of Economics*, **125** (1), 1–45.
- COLE, S., FERNANDO, A. N., STEIN, D. and TOBACMAN, J. (2016). Field comparisons of incentive-compatible preference elicitation techniques.
- COLE, S. A. and FERNANDO, A. N. (2012). The value of advice: Evidence from mobile phone-based agricultural extension.
- DRACA, M., SADUN, R. and VAN REENEN, J. (2006). Productivity and ict: A review of the evidence.

- DUFLO, E., KREMER, M. and ROBINSON, J. (2011). How high are rates of return to fertilizer? evidence from field experiments in kenya. *Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence in Kenya*.
- FEDER, G. (2005). The challenges facing agricultural extension - and a new opportunity. *New Agriculturalist*.
- , LAU, L. and SLADE, R. (1987). Does agricultural extension pay? the training and visit system in northwest india. *American Journal of Agricultural Economics*, **69** (3).
- FERNANDO, A. N. (2016). Social interactions, technology adoption and information exchange: Evidence from a field experiment.
- FOSTER, A. and ROSENZWEIG, M. (2010). Microeconomics of technology adoption. *Annual Reviews of Economics*.
- GANDHI, R., VEERARAGHAVAN, R., TOYAMA, K. and RAMPRASAD, V. (2009). Digital green: Participatory video and instruction for agricultural extension. *Information Technologies and International Development*, **5** (1).
- GLENDENNING, C., BABU, S. and ASENSO-OKYERE, K. (2010). Review of agricultural extension in india: Are farmer' information needs being met? *IFPRI Discussion Paper 01048*.
- GROUP, W. B. (2012). *World Development Indicators 2014*. World Bank Publications.
- HSIEH, C.-T. and KLENOW, P. J. (2009). Misallocation and manufacturing tfp in china and india. *The Quarterly Journal of Economics*, **124** (4), 1403–1448.
- INDIA, T. R. A. (2015). Annual report 2014-2015.
- JACK, K. (2011). Market inefficiencies and the adoption of agricultural technologies in developing countries. *ATAI*.
- KLING, J. R., LIEBMAN, J. B. and KATZ, L. F. (2007). Experimental analysis of neighborhood effects. *Econometrica*, **75** (1), 83–119.
- MCKENZIE, D. and WOODRUFF, C. (2012). What are we learning from business training and entrepreneurship evaluations around the developing world? *World Bank Policy Research Working Paper*, (6202).

PATEL, N., CHITTAMURU, D., JAIN, A., DAVE, P. and PARIKH, T. (2010). Avaaaj otalo - a field study of an interactive voice forum for small farmers in rural india. *Proceedings of ACM Conference on Human Factors in Computing Systems*.

SAITO, K. and WEIDEMANN, C. (1990). Agricultural extension for women farmers in africa. *World Bank*.