

‘MOBILE’ISING AGRICULTURAL ADVICE: TECHNOLOGY ADOPTION, DIFFUSION AND SUSTAINABILITY*

Shawn A. Cole and A. Nilesh Fernando

Mobile phones promise to bring the ICT revolution to previously unconnected populations. A two-year study evaluates an innovative voice-based ICT advisory service for smallholder cotton farmers in India, demonstrating significant demand for, and trust in, new information. Farmers substantially alter their sources of information and consistently adopt inputs for cotton farming recommended by the service. Willingness to pay is, on average, less than the per-farmer cost of operating the service for our study, but likely exceeds the cost at scale. We do not find systematic evidence of gains in yields or profitability, suggesting the need for further research.

Differences in technology adoption drive productivity differences in agriculture. In turn, a variety of observers have pointed out that access to information and awareness of agricultural technologies may play an important role in their adoption (Jack, 2011). Yet in-person extension services are expensive, slow and cumbersome. In India, dispersed rural populations, monitoring difficulties, and limited accountability severely constrain the reach of in-person extension systems: fewer than 6% of the agricultural population reports having received information from these services.¹

The rapid spread of mobile phones holds the promise of bringing high-quality advice to billions of previously unconnected individuals. Yet, we know relatively little about whether individuals with low levels of technological literacy will trust information and whether the provision of information by phone will change behaviour. This paper examines whether the introduction of an information service that is able to deliver timely, relevant and actionable information to farmers can meaningfully influence agricultural practices. Specifically, we evaluate Avaaj Otalo (AO), a mobile phone-based technology service that both pushes information to farmers via voice calls, and allows users to call a hotline, ask questions and receive a recorded response from agricultural

* Corresponding author: A. Nilesh Fernando, Department of Economics, University of Notre Dame, Notre Dame, IN 46556, USA. Email: nilesh.fernando@nd.edu

This paper was received on 31 May 2019 and accepted on 24 June 2020. The Editor was Ekaterina Zhuravskaya.

The data and codes for this paper are available on the Journal website. They were checked for their ability to reproduce the results presented in the paper.

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

This study was registered in the AEA RCT registry as AEARCTR-0000678. 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 organisation, Precision Agriculture for Development, seeking to further test, and bring to scale, mobile phone-based agricultural advice services; this non-profit funded a postdoctoral 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 developing and 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, HK Seo and Chungun Yoon 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 numerous colleagues and participants at a variety of seminars.

¹ This 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 these data.

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 400 households (hereafter, 'AO group') and to test the hypothesis that ICT would not be effective without at least some in-person element, an additional 400 households received both AO and an annual in-person extension session (hereafter, 'AOE group'). A further 400 households served as a pure control group. The households were spread across 40 villages in Surendranagar district in Gujarat, India, and randomisation 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. An important difference from prior ICT-based agricultural extension programs is that the information was exclusively delivered through voice messages as opposed to text-based approaches that may be less suited to semi-literate environments. 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 were asked, prior to the intervention, to identify the individuals with whom they discussed farming.

We find considerable demand for the service: nearly 90% of the combined treatment group (AO+AOE) called into the AO line over two years and 40% asked a question. The average treatment respondent used the service for almost seven hours (median 5.2 hours), making 22 calls, calling into the service for more than 2.5 hours and listening to four additional hours of push content. In addition, we find that respondents in the AOE group used the service for an hour more than the AO group (7.3 hours to 6.2 hours). The service increased subjective trust in mobile phone-based programs as a source of agricultural advice, from nearly 0 points at baseline to more than 6 points on a 10-point scale by the endline.

Aggregating across survey rounds and treatments, we find a 5.8 standard deviation increase in the reported use of mobile phone-based information in agricultural decision-making, as measured by an index aggregating a variety of input use and sowing decisions. Farmers relied less on commissions-motivated agricultural input dealers for pesticide advice, and less on their prior experience for fertiliser-related decisions. Our outcomes are largely self-reported, raising the concern of desirability bias. Our survey teams were distinct from the service team (and not affiliated with the NGO). We are further reassured by the fact that farmers in treatment groups did not report any significant changes in the use of mobile phone-based information for price information—which the service did not provide; and that self-reported and server logged usage of the service are virtually identical.²

Relative to in-person extension programmes (Bardhan and Mookherjee, 2011; Duflo *et al.*, 2011) and interventions in markets linked to agricultural productivity (Giné and Yang, 2009; Cole *et al.*, 2013) that often involve costly on-site infrastructure and labour resources, our intervention represents a potential low-cost way to promote technology adoption. However, even with considerable usage, the largely on-demand nature of the AO service means that respondents are not provided a homogeneous treatment, but rather recommendations tailored to their specific needs. This poses a challenge to evaluation, as it is difficult to specify *ex ante* what one should expect the effect of the service to be, and because different farmers receive different messages.

² We acknowledge, however, that this still leaves open the possibility of sophisticated demand effects where farmers over-emphasise the usefulness of features of AO.

We address this by testing whether a broad set of agricultural practices, aggregated by crop type and input type, respond to the treatment.

First, at the crop level, farmers who received access to the AO service were significantly more likely to adopt recommended agricultural inputs for cotton cultivation—their primary crop—as measured by an aggregate index (0.13 SD). Looking at indices aggregating types of inputs, we find that the treatment increased the adoption of recommended seeds by 0.09 SD. In addition, accounting for spillovers these estimates increase to 0.15 SD and 0.11 SD for the cotton and seed indices, respectively. Importantly, the intervention induced consistently higher expenditure on irrigation (15% relative to the baseline control mean) and seeds (20% relative to baseline control mean), which we interpret as a complementary investment to more input-intensive cultivation practices which were recommended by the service. We find some evidence to suggest that the intervention influenced agricultural knowledge as measured by 44 questions (3% overall and 5% for cotton-related questions, relative to the baseline control mean). We suggest some caution in interpreting the estimates for inputs and knowledge: while the per-comparison *p*-values are significant (and relevant), adjusting these *p*-values for testing multiple hypotheses leaves them well above standard thresholds for rejection.

The primary goal of our experiment was to evaluate whether the service could facilitate changes in technology adoption by farmers; in this we find the program successful. We do not, however, find systematic evidence that the intervention increased crop yields or profit. While the point estimates are often positive, they are noisy, as smallholder yields, and particularly self-reported yields, are quite noisy (Lobell *et al.*, 2020). In addition, we note that even where farmers follow practices that are beneficial in expectation, the stochastic nature of rainfall complicates the detection of treatment effects (Rosenzweig and Udry, 2020). *Ex post*, our single study is therefore likely underpowered to capture effects on these outcomes; we believe this is a crucial area for further study, including meta-analyses.

Our treatment also induces variation in the availability of information in social networks, allowing us to estimate peer effects. Here, the evidence is mixed. We find some evidence for complementarities among treated respondents: treated respondents with more treated peers (social network members) are more likely to call into the service and use the service for longer. In addition, among a set of respondents who are peers of study respondents, we find that exposure to a treated respondent results in lower pest-related cotton losses. On the other hand, we don't find that study respondents—either in treatment or control—exposed to the treatment through a peer, are more likely to change the sources of information they use in agricultural decision-making.

At the time of the endline, we conducted a series of willingness to pay (WTP) experiments to estimate demand for AO. Average WTP for a nine-month AO subscription across multiple price elicitation methods is roughly \$2, compared to a cost of provision for the same period of \$7. While suggestive of low WTP, it is worth noting here that as this service scales, the per-farmer cost of provision reduces substantially as the marginal cost of scaling is very low.³

A large literature focuses on the microeconomics of technology adoption (for a survey, see Foster and Rosenzweig, 2010). We contribute to this literature by examining whether an information service can facilitate improved production practices. Our work complements a developing literature on the potential for digital agricultural extension (see Fabregas *et al.*, 2019 for a recent review). Prior experimental work has used SMS messages to send farmers agricultural advice:

³ The non-profit organisation Precision Agriculture for Development (Cole is on the board of this organisation) estimates the current cost of approximately \$2 per farmer per year of a similar service serving almost 600,000 farmers in Odisha, India.

Fabregas *et al.* (2020) find that these messages increased the take-up of agricultural lime in Kenya and Rwanda, Casaburi *et al.* (2020) find mixed evidence on sugarcane yields in Kenya, while Fafchamps and Minten (2012) do not find that SMS messages influenced cultivation practices in India. Among maize farmers in Uganda, Van Campenhout (2017) find that extension videos influenced cropping patterns. Our treatments differ from much of the previous work in that participants receive voice messages as opposed to SMS messages or videos. In addition, the flow of information is demand-driven and customised according to the needs of individual farmers rather than aggregated at the level of crop choice or region.

More generally, this paper advances the literature on the efficacy of agricultural extension (Feder *et al.*, 1987; Gandhi *et al.*, 2009; Duflo *et al.*, 2008). 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 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. BenYishay and Mobarak (2019) compare the impact of incentivised extension agents to non-incentivised extension agents in Malawi, finding that incentives affect extension effort and that the identity of the extension agent affects the adoption of information.

We demonstrate that informational inefficiencies are real and that farmers are aware they lack information: there is considerable demand for high-quality agricultural information.⁴ Our results complement recent work that measures productivity enhancement from ICTs in developed countries (Draca *et al.*, 2007).

Finally, we carefully evaluate how selling the service, rather than giving it away for free, would impact access to the service. We provide an experimental comparison of the Becker–DeGroot–Marschak (BDM) mechanism to standard sales offers, demonstrating that the BDM mechanism, which requires a smaller number of data points, yields credible estimates of the demand curve.⁵

This paper is organised as follows. The first section provides context and the details of the AO intervention. Section 2 presents the experimental design and the empirical strategy, while Section 3 presents the results from the two years of survey data. Following this, Section 4 considers threats to the validity of the results, and Section 5 concludes.

1. Context and Intervention Description

1.1. Agricultural Extension

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 access to information and awareness of agricultural technologies may play an important role (Jack, 2011).

According to the World Bank, there are more than one million agricultural extension workers in developing countries, and public agencies have spent over \$10 billion on public extension programs in the past five decades (Feder, 2005). The in-person extension model, ‘Training and

⁴ Informational 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 utilised (Jack, 2011).

⁵ In a companion paper, (Cole *et al.*, 2020) we explore these findings in more detail.

Visit' extension, has been promoted by the World Bank throughout the developing world and is generally characterised 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 centres, 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 according to the Reserve Bank of India.

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 workforce of public extension agents. However, evidence of the efficacy of these extension services is limited. In India, dispersed rural populations, monitoring difficulties, and a lack of accountability hamper the efficacy of in-person extension systems: fewer than 6% of the agricultural population reports having received information from these services.⁶

Yet, in-person 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 capacity: 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 on marginalised 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 may be difficult to overstate (Saito and Weidemann, 1990; Birkhaeuser *et al.*, 1991). 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

⁶ See footnote 1.

or face incentives to recommend the wrong product or excessive dosage (Anderson and Birner, 2007).⁷

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 have opened up the possibility of using a completely different model to deliver agricultural extension services.

1.2. *Avaaj Otalo: Mobile Phone-Based Extension*

Roughly 50% of the Indian labour 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 nationally.⁸ Mobile phone access has fundamentally 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 utilises an innovative information technology service, Avaaj 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 is 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 customised 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 in-person 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 in-person 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

⁷ An audit study we conducted 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. Our findings suggest that the information provided is rarely customised to the 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. See Cole and Sharma (2020) for further details.

⁸ These figures are calculated using the Annual Report of the Telecom Regulatory Authority of India (TRAI, 2015) and the World Bank Development Indicators (World Bank 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 out of a total workforce of 497 million.

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 minimise 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 the 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 utilisation of the information.

However, while AO overcomes many of the challenges of in-person 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 in-person extension administered through staff at DSC and toll-free access to AO.

2. Experimental Design and Empirical Strategy

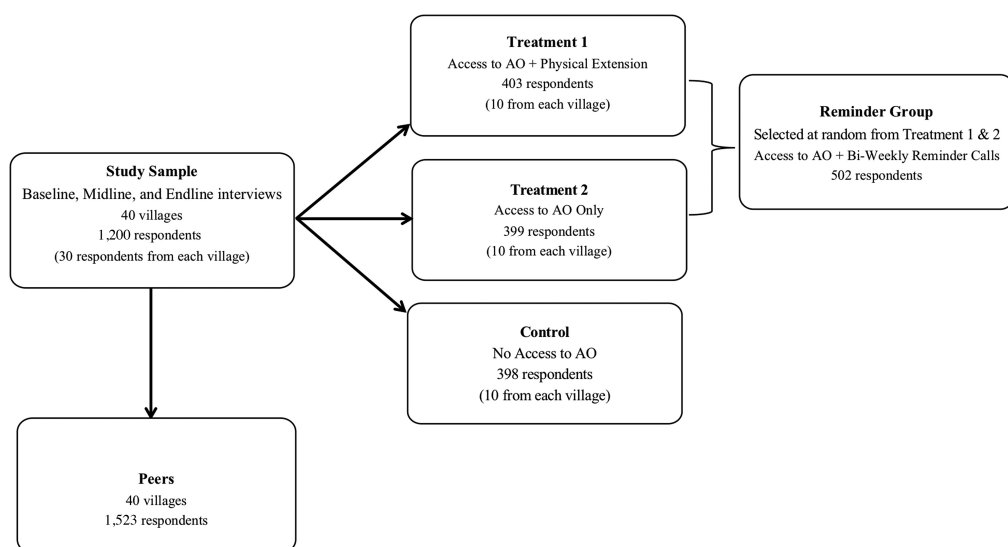
We selected two administrative blocks,⁹ Chotila and Sayla, in the Surendranagar district of Gujarat as the site of the study, as our field partner, DSC, had done work in the area. Farmer lists, consisting of all households that grew cotton and owned a mobile phone, were created in 40 villages and served as our sampling frame.

We invited randomly selected farmers from this set to participate in a study (farmers were not told that the study involved mobile phones, nor were they told treatment status, when agreeing to participate in the study). Nearly all agreed to participate, and we obtained a sample of 1,200 respondents, 30 from each village.¹⁰ Figure 1 summarises 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 in-person extension. The in-person extension component consisted of a single session each year lasting roughly two-and-a-half hours on DSC premises in Surendranagar.¹¹ The second treatment group (hereafter, AO) received toll-free access to AO, but no offer of in-person agricultural extension, and the final set of households served as the

⁹ A block is an administrative unit below the district level.

¹⁰ Online Appendix Table A5 provides further details on sample selection.

¹¹ The in-person agricultural extension program was, in some sense, rather 'light touch', involving only a single meeting between the farmer and the extension team each year. This treatment arm was meant to address a potential concern about the AO service, that farmers would not trust a purely digital intervention.

Fig. 1. *Experimental Design.*

control group. In addition, among the two treatment groups (AO and AOE), 500 were randomly selected to receive biweekly 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 2011, and a phone survey consisting of 798 respondents was completed in November, 2011.¹² The midline survey was completed by August 2012, and the endline survey was completed by August 2013.

To gauge balance, 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 comparing the treatment to the control mean. We use the ANCOVA specification as suggested by McKenzie (2012) in order to increase our power to detect effects, given the low autocorrelation of most outcomes in our data. Specifically, our main specification only uses the midline and endline data and controls for the baseline value of the outcome of interest:

$$y_{ivt} = \alpha_v + \alpha_t + \beta_1 \text{Treat}_{iv} + \beta_2 y_{iv0} + \varepsilon_i, \quad (2)$$

where α_v and Treat_{iv} are as above, α_t is a fixed effect for the survey round and y_{iv0} is the baseline value of the outcome of interest.

¹² The previous version of this paper (Cole and Fernando, 2012) analysed treatment effects using results from this phone survey.

Date	Event
May/2011	Cotton planting decisions begin
May/2011	Listing for baseline survey
Jul/2011	Baseline (Paper) Survey
Aug/2011	AO training for treatment respondents
Aug/2011	AO service activated for all treatment respondents
Sep/2011	Reminder calls started
Nov/2011	Physical Extension Round 1
Nov/2011	Phone Survey Round 1
Dec/2011	Phone Survey Round 2
Mar/2012	Peer Survey
Jun/2012	Midline (Paper) Survey
Aug/2012	AO training for treatment respondents Round 2
Oct/2012	Field visits to gather information on Rabi planting decisions
Nov/2012	Peer Survey Part 2
Nov/2012	Physical Extension Round 2
Mar/2013	Phone Survey 3
Jul/2013	Endline (Paper) Survey
Jul/2013	Willingness to Pay Study
Jul/2013	Ending push calls/intervention

Fig. 2. *Project Timeline.*

While increasing statistical power, the decision to randomise at the household rather than village level raises the possibility that the control group may also have access to information through our treatment group. This suggests that any treatment effects may in fact underestimate the value of the service.

In order to systematically assess this concern, we include a specification—denoted as ‘spillover’—in all the main tables that controls for potential spillovers. Specifically, at baseline, we asked all respondents to list the three contacts with whom they most frequently discussed agricultural information. As such, we are able to identify whether any of these peers also received the treatment. We amend the ANCOVA specification above to include a set of fixed effects for the number of peers listed at baseline and a control for the fraction of these peers who received the treatment as below:

$$y_{ivt} = \alpha_v + \alpha_t + \beta_1 \textit{Treat}_{iv} + \beta_2 y_{iv0} + \beta_3 \textit{Treat_Frac}_{iv} + \sum I(\# \textit{Peers} = i)_{iv} + \varepsilon_i, \tag{3}$$

where α_v and α_t are as above, $\sum I(\# \textit{Peers} = i)_{iv}$ is a fixed effect for the number of peers listed by a respondent at baseline and $\textit{Treat_Frac}_{iv}$ is the fraction of these peers who are assigned to treatment.

While we view the above specification as a ‘control’ that allows us to compare how the treatment effect β_1 changes with and without spillover controls, we also view spillover effects as being of independent interest.

As such, in Table 6 we estimate the above specification for all study respondents and we separately estimate the effect of being exposed to a treated respondent for non-study respondents. In particular, using the 'peer survey' we collected information on 1,114 non-study respondents, i.e., peers listed by study respondents at baseline who were not themselves a part of the study.¹³ We estimate the extent of such peer effects or information spillovers with the following specification:

$$y_{iv} = \alpha_v + \beta(Treat\ References/References)_{iv} + \sum_{i=2}^7 I(\# References = i)_{iv} + \varepsilon_{iv}, \quad (4)$$

where α_v is as above, $\sum_{i=2}^7 I(\# References = i)_{iv}$ is a fixed effect for the number of study respondents who list a peer as a top agricultural contact and $(Treat\ References/References)_{iv}$ 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 it was not always clear *ex ante* which subjects farmers would inquire about. We address concerns about multiple inference in four ways. First, we use the content generated by farmers, and by our agronomist, as a broad guide for conducting empirical analysis.¹⁴ Second, we aggregate agricultural practices into indices, following, for example, Kling *et al.* (2007). Our agronomist characterised all agricultural practices reported in the survey as either consistent with best practices (we assign 1 to these responses), or either inconsistent or indeterminate (we assign a 0 to these responses). We then aggregate all variables corresponding to recommended practices by calculating a z-score for each component and take the average z-score across components. It is important to note that this is not a quantity index (i.e., more pesticide or fertiliser does not increase the corresponding z-score in a deterministic manner). Rather, components of the z-score are positive for appropriate use of inputs. Each component z-score is computed relative to the control group mean and standard deviation at baseline. The components of the index are weighted by the inverse of the covariance matrix to adjust for highly correlated outcomes (Anderson, 2008).

Third, we address the importance of Type 1 error directly in two ways. First, in Online Appendix Table A2 we report the number of comparisons that are statistically different from zero at conventional levels of statistical significance in each survey round and for each treatment arm. Panel A shows that at baseline, the number of comparisons that are significantly different from zero is consistent with what we would expect given Type 1 error at each level of significance. In contrast, the analogous results for the midline (Panel B) and endline (Panel C) reveal that we are able to reject the null for a much larger number of the same comparisons than would be predicted by Type 1 error.

Finally, we use both the standard Bonferroni–Holm and the Westfall–Young correction to compute a family-wise error rate (FWER) across our main outcomes of interest (See Online Appendix Table A9). The latter correction uses randomisation inference to compute a family-wise error rate (FWER) for a set of comparisons. By re-estimating the full set of outcomes this correction takes into account the correlation between outcomes and is less conservative than the Bonferroni–Holm method. We separately compute a FWER for input adoption (Panel B) and for an overall set of summary indices (Panel A) (Anderson, 2008).

¹³ As we note in Figure 1, the peer survey included 1,523 respondents, 409 of whom were study respondents and the remaining 1,114 were non-study respondents. Of those who were study respondents, 143 belonged to the control group, 140 belonged to the AOE group, and 126 belonged to the AO group.

¹⁴ See Online Appendix Table A1 for details of questions asked by farmers on the AO service and push content provided.

2.1. Summary Statistics and Balance

In this section, we assess balance between the treatment group that received access to the advisory service and in-person extension ('AOE'), the treatment group that only received access to the advisory service ('AO'), the combination of these two treatment groups ('Treat') and the control group.

Table 1 contains summary statistics for age, education, profit from agriculture, 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 and column (2) tests the initial randomisation balance between the combined treatment group and the control group. Column (3) tests the initial randomisation balance between the AOE group and the control group, column (4) tests the balance between the AO group and the control group and column (5) tests for balance between the AOE group and AO group.

We see that respondents are on average 46 years old and have approximately four years of education. Columns (3)–(5) show that the randomisation was largely successful for the treatment groups across demographic characteristics (Panel A) and indices capturing information sources, crop-specific and general input use (Panel B). However, an imbalance exists in the index for wheat management between the AOE group and control and another imbalance exists in the area of cotton planted between the AO group and the control group.¹⁵ However, the latter imbalance exists in 2010 but not in 2011 (both periods are prior to treatment).¹⁶ Both treatment groups are also more likely to grow wheat, but this crop is mostly grown for home consumption in this context.

Particularly as cotton is the most important crop in our sample, we understand the importance of systematically accounting for baseline differences in covariates across treatment groups. As such, in the tables that follow we adopt the double LASSO machine learning approach (DML) to pick an optimal set of control variables as proposed by Belloni *et al.* (2014). Online Appendix Table A18 details the set of control variables (including their interactions) to which we apply the algorithm.¹⁷

3. Experimental Results

In the results we detail below, we report estimates comparing the combined treatment group ('Treat') to the control group—hereafter, the 'treatment group'—rather than its constituent treatment arms (the 'AOE' group and the 'AO' group).¹⁸ Similarly, we do not present the effects of the reminder treatment in the main tables. In both cases, this is to streamline the discussion and presentation of our results because the treatment effects for the AOE group and the AO group rarely differ, in addition to there rarely being a marginal effect of the reminder treatment over and

¹⁵ Online Appendix Table A20 details the variables used to construct all aggregate indices.

¹⁶ Note, the 2011 figures for wheat and cumin are not reported, as they are grown during the Rabi season after the treatment was administered.

¹⁷ Additionally, Online Appendix Table A2 provides a more systematic treatment of balance in our sample. We look for significant differences in baseline characteristics between the treatment groups and control respondents. Among the differences computed using the latter specification (examining all 1,643 baseline variables), we find that 0.7% are significantly different from zero at the 1% level, 4.3% are different at the 5% level of significance and 8.6% at the 10% level. These results suggest that the randomisation was successful and that the imbalances are a result of chance rather than any systematic mistake in the randomisation mechanism.

¹⁸ We report the disaggregated results by both the AOE and AO treatments as well as the reminder treatment in Online Appendix Table A12.

Table 1. *Summary Statistics and Balance.*

Dependent variable	Control mean (Baseline) (1)	Treat-control ITT (2)	AOE-control ITT (3)	AO-control ITT (4)	AOE-AO ITT (5)
<i>Panel A: Demographic characteristics</i>					
Age	46.539 [15.161]	-0.369 (0.915)	-0.811 (1.047)	0.023 (1.061)	-0.844 (1.049)
Years of education	4.235 [3.836]	-0.187 (0.230)	-0.116 (0.263)	-0.266 (0.273)	0.141 (0.272)
Landholdings—acres	6.077 [5.596]	0.095 (0.332)	0.236 (0.385)	-0.064 (0.393)	0.368 (0.407)
Profit from agriculture (Rupees, winsorised fraction = 0.01)	1.36e+05 [1.26e+05]	5,082.579 (7,665.933)	3,972.733 (8,766.187)	6,184.739 (9,124.237)	-1,552.060 (9,124.480)
<i>Panel B: Indices (SD units)</i>					
Mobile phone-based Information usage	0.000 [1.000]	0.055 (0.071)	0.012 (0.077)	0.096 (0.096)	-0.076 (0.097)
Cotton management	0.000 [1.000]	0.001 (0.059)	-0.036 (0.067)	0.036 (0.070)	-0.081 (0.071)
Wheat management	0.000 [1.000]	-0.070 (0.057)	-0.126** (0.059)	-0.010 (0.069)	-0.116** (0.059)
Cumin management	0.000 [1.000]	-0.053 (0.056)	-0.052 (0.063)	-0.051 (0.063)	0.003 (0.060)
Pesticide management	0.000 [1.000]	-0.081 (0.057)	-0.063 (0.066)	-0.097 (0.066)	0.035 (0.067)
Fertiliser management	0.000 [1.000]	-0.027 (0.059)	-0.061 (0.068)	0.006 (0.070)	-0.069 (0.068)
<i>Panel C: Agricultural activity</i>					
Planted cotton (2010)	0.985 [0.122]	-0.003 (0.008)	0.002 (0.008)	-0.008 (0.010)	0.008 (0.009)
Area cotton planted (2010) (acres)	4.448 [3.622]	0.422* (0.232)	0.278 (0.281)	0.575** (0.269)	-0.271 (0.293)
Area cotton planted (2011) (acres)	4.990 [3.846]	0.293 (0.247)	0.384 (0.309)	0.209 (0.280)	0.217 (0.318)
Planted wheat (2010)	0.776 [0.417]	-0.053** (0.025)	-0.053* (0.029)	-0.053* (0.029)	-0.005 (0.030)
Area wheat planted (2010) (acres)	1.171 [1.346]	0.016 (0.089)	-0.101 (0.088)	0.134 (0.122)	-0.226** (0.114)
Planted cumin (2010)	0.425 [0.495]	-0.018 (0.028)	-0.013 (0.033)	-0.024 (0.032)	0.011 (0.032)
Area cumin planted (2010) (acres)	0.762 [1.406]	-0.019 (0.083)	-0.055 (0.095)	0.014 (0.097)	-0.059 (0.096)
N	398	1,200	801	797	802

Notes: This table reports summary statistics and assesses balance across groups using data from the Baseline survey, conducted between June 26 and August 11, 2011. Participants were randomised into three groups. AO group received AO access. AOE group received AO access and physical extension. 'Treat' refers to the combined treatment group. The control group received neither treatment. The indices aggregate information over multiple outcomes for which we expect unidirectional treatment effects. Each index consists of the average of the z-scores for each component of the index, with the control group mean and standard deviation as reference. Mobile phone-based information usage index: aggregates mobile phone use across crop decision, soil preparation, pest identification, weather, cotton pesticides, cotton fertilisers, wheat fertilisers, cumin pesticides, and cumin fertilisers. Management practices indices: seed usage + pesticide purchase + pesticide usage + fertiliser purchase + fertiliser usage for the three different crops—cotton, wheat, and cumin. Pesticide management index: dummy to indicate purchase/use of a pesticide. Fertiliser management index: dummy to indicate purchase/use of a fertiliser. Seed management index: dummy to indicate purchase/use of recommended seeds (see Online Appendix Table A20 for index details). Profit from agriculture refers to the difference between total income from of all crops grown less total input expenditure in the past year, where the latter includes seeds, fertilisers, irrigation, pesticides, hired labour, and household labour priced at the mean wage of hired labour. Column 1 shows the summary statistics (mean and standard deviation) for the control group at baseline. Columns 2–4 report an Intention to Treat (ITT) estimate of the difference in means (and robust standard error) between the treatment group specified and the control group. Column 5 reports the difference in baseline covariates between the AO and AOE group. All specifications include village fixed effects. Asterisks denote statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2. *Usage of the AO Information Service.*

Dependent variable	Control mean (midline) (1)	Treat–control (midline) (2)	Control mean (endline) (3)	Treat–control (endline) (4)
Incoming AO usage + push calls listened (minutes)	0.000 [0.000]	214.857*** (11.078)	0.168 [1.937]	421.211*** (19.681)
Called in to the AO line	0.000 [0.000]	0.660*** (0.018)	0.003 [0.053]	0.888*** (0.012)
Total number of calls	0.000 [0.000]	8.845*** (1.069)	0.006 [0.106]	22.861*** (2.577)
Total incoming AO usage (minutes)	0.000 [0.000]	93.924*** (10.743)	0.002 [0.038]	167.811*** (18.904)
Asked a question	0.000 [0.000]	0.313*** (0.018)	0.008 [0.091]	0.401*** (0.019)
Percentage of total push call time listened to	0.000 [0.000]	0.550*** (0.007)	0.000 [0.004]	0.551*** (0.008)
N	398	1,123	359	1,080

Notes: This table reports usage statistics collected on the AO server. ‘Treat’ group refers to the 802 farmers that received access to AO. Column 1 provides the mean and standard deviation for the control group at midline. Column 2 reports the treatment effect estimate from a simple difference specification at midline. Column 3 provides the mean and standard deviation for the control group at endline. Column 4 reports the treatment effect estimate from a simple difference specification at endline. All specifications include village fixed effects. Asterisks denote statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

above that of the combined treatment group. However, there are a few exceptions to this general characterisation which we discuss below.

We acknowledge that this means we estimate a composite treatment effect that includes a weighted-average of the treatment groups.¹⁹ We estimate treatment effects averaged across survey rounds (as in equation (2)) rather than show effects by round, aside from cases where the evolution of treatment effects (e.g., adoption of the service in Table 2) or the stochastic nature of results is especially salient (e.g., yield effects are dis-aggregated in Online Appendix Table A17).

In the tables that follow, unless otherwise noted, we report estimates from the ANCOVA specification from equation (2), the spillover specification from (3), and a simple difference specification (1) using the double LASSO procedure (hereafter, DML) to select control covariates.²⁰

3.1. *Take-up and Usage of AO*

Table 2 reports information on take-up and usage of the AO service at midline and endline.²¹ While control respondents were not barred from AO usage, we did not inform them of the service, and those who did use it had to pay their own airtime costs. 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. Note, as there is no baseline usage of the service the estimates here are simple differences rather than from the ANCOVA specification.

¹⁹ As budgetary restrictions prevented a full 2×2 factorial design, we are unable to implement the corrective methods proposed by Muralidharan *et al.* (2020). Moreover, we note that when we estimate the ‘long’ model, the qualitative conclusions of our analysis rarely change.

²⁰ Note, in the DML specification the algorithm chooses whether to include the baseline value of the outcome as a control.

²¹ Note, we disaggregate results by midline and endline here to show the evolution of usage.

Driven primarily by the 'push' features, adoption of the service was broad and deep. By the endline, nearly 90% of treatment respondents had called into the service, and the mean total usage for the combined treatment group—including both incoming calls and time spent listening to push calls—was 7 hours (median 5.2 hours). Note, in this case we do see a significant difference of roughly 1 hour between usage of respondents in the AOE group and the AO group (see Online Appendix Table A12).

Overall, we find that inbound usage was more concentrated: average time spent calling in by the endline was 2.7 hours with 22 calls, however the top 10% of users accounted for 70% of incoming call time. In contrast, interest in information sent through push calls was more evenly skewed: treatment respondents listened to 55% (median 58%) of total push call content by endline.²² By the endline, 40% of treated respondents had asked a question about their agriculture on the system.²³

Taken together, these results represent substantial induced usage for treatment farmers. Additionally, these effects also mask important temporal patterns shown in Figure 3, which reports average service 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 users acquiring a stock of knowledge and supplementing thereafter with dynamic information needs throughout the season.

In addition to providing a temporally relevant flow of knowledge, 40% of the combined treatment group received customised answers to agricultural questions. Online Appendix Table A1 provides a categorisation 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, 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. A larger study might have experimentally varied the topics of the push content (for example, either matching or not matching subjects of contemporaneous queries), to further disentangle the role of push vs. pull information. However, operational and power considerations precluded inducing such variation in this study.


3.2. *Sources of Information and Agricultural Knowledge*

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. On average, treatment farmers are 66 percentage points (p.p.) more likely to report using mobile phone-based information to make agricultural decisions. The treatment effects on the reported level of trust in mobile phone-based information are also substantially higher: approximately 5.8 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


²² This is calculated as follows: the total minutes of content listened to by all users, divided by the total minutes of push content that would have been heard had all users listened to the entirety of each message.

²³ In Online Appendix Table A12, we also find that the reminder treatment significantly increased service usage: the reminder group had used the service almost an hour more on average by midline and over 90 minutes more by endline, but were not statistically more likely to call into the line.

Panel A: Average AO Usage by Month



Panel B: Average Percentage of Push Call Listened to By Month



List of experiment/activity by month:

- | | |
|--|--|
| 09 - (Sep 2011) AO Push Calls Start | 19 - (Jul 2012) Midline Survey Ends |
| 09 - (Sep 2011) Reminder Calls Start | 22 - (Oct 2012) Cotton Harvest Starts (Year 2) |
| 10 - (Oct 2011) Physical Extension (Round 1) | 23 - (Nov 2012) Physical Extension (Round 2) |
| 11 - (Nov 2011) Cotton Harvest Starts (Year 1) | 30 - (Jun 2013) Cotton Sowing Starts (Year 3) |
| 18 - (Jun 2012) Cotton Sowing Starts (Year 2) | 31 - (Jul 2013) Endline Survey Starts |
| 18 - (Jun 2012) Midline Survey Starts | 32 - (Jul 2013) Endline Survey Ends |

Fig. 3. Use of Mobile Phone-Based Information by Month.

Notes: This figure reports monthly usage of the AO service (both incoming and outgoing) based on data collected from the AO server. Panel A represent total incoming AO usage (using the missed call service) to either record a question, listen to other messages or share experiences with other users. Panel B represents the percentage of total push calls (outgoing weekly agricultural messages) listened to. The service was provided to all treatment farmers from September 2011 to August 2013. ‘AO+AOE’ group refers to the 802 farmers that received access to AO. AOE group includes 403 farmers who had access to AO and physical extension. AO group refers to the 399 farmers who only had access to AO. ‘Reminder’ calls refer to the 502 farmers from the ‘AO+AOE’ group that received biweekly calls encouraging them to use the service.

Table 3. *Effects on Sources of Agricultural Information and Knowledge.*

Dependent variable	Control mean (baseline) (1)	Treat–control ANCOVA (2)	Treat–control spillover (3)	Treat–control DML (4)
<i>Panel A: Across all agricultural decisions</i>				
Index of mobile phone-based information usage (SD units)	0.000 [1.000]	5.543*** (0.246)	5.678*** (0.252)	5.606*** (0.246)
Used mobile phone-based information	0.093 [0.291]	0.665*** (0.017)	0.671*** (0.017)	0.661*** (0.017)
Trust in mobile phone-based information (on a scale of 1–10)	0.606 [2.031]	5.882*** (0.151)	5.958*** (0.149)	5.890*** (0.151)
N (Trust)	398	2194	2194	2194
<i>Panel B: Usage of mobile phone-based information by decision type</i>				
Crop decision	0.000 [0.000]	0.050*** (0.006)	0.052*** (0.006)	0.049*** (0.006)
Pest management	0.000 [0.000]	0.168*** (0.011)	0.172*** (0.011)	0.168*** (0.010)
Fertiliser management	0.003 [0.050]	0.086*** (0.008)	0.088*** (0.008)	0.085*** (0.008)
Weather	0.003 [0.050]	0.291*** (0.014)	0.298*** (0.014)	0.294*** (0.014)
Soil preparation	0.000 [0.000]	0.030*** (0.005)	0.032*** (0.005)	0.030*** (0.005)
Prices	0.023 [0.149]	0.002 (0.009)	0.004 (0.009)	0.005 (0.009)
<i>Panel C: Agricultural knowledge</i>				
Total correct answers to questions (44 questions)	14.156 [5.279]	0.350 (0.232)	0.426* (0.234)	0.411* (0.222)
Cotton-related (20 questions)	4.774 [2.061]	0.199 (0.132)	0.262** (0.133)	0.251** (0.126)
N	398	2,203	2,203	2,203

Notes: This table reports the impact of AO on WTP, agricultural productivity, and profit. The results use data from both the Midline survey and the Endline survey. Input expenditure includes total money spent on seeds, fertilisers, irrigation, pesticides, hired labour, and household labour priced at the mean wage of hired labour. Profit from agriculture refers to the difference between total income from all crops grown less total input expenditure in the past year. WTP reports the bid in a BDM demand elicitation game in Rupees. ‘Treat’ group refers to the 802 farmers that received access to AO. Column 1 provides the mean and standard deviation for the control group at baseline. Column 2 reports the treatment effect estimate from an ANCOVA specification. Column 3 uses the ANCOVA specification and controls for the baseline treatment status of a respondent’s peer group to assess spillover effects. Column 4 reports the ANCOVA specification and uses double ML to pick an optimal set of control variables. All specifications include village fixed effects and survey round fixed effects. Asterisks denote statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

subject areas is 5.5 standard deviations higher across the treatment group and is significant at the 1% level.

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 treated respondents consistently report using mobile phone-based information across a series of agricultural decisions. In particular, we observe large effect sizes in the case of pest management (17 p.p.) and smaller effects in the case of fertiliser decisions (9 p.p.) and crop planning (5 p.p.).

Other than input-related decisions, mobile phone information is also used by the treatment group for other topics such as weather information (30%). 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. Accounting for spillover effects consistently, though marginally, increases coefficient estimates. Likewise, using DML to pick control covariates leaves both the point estimates and the precision with which they are estimates largely unchanged.²⁴

Next, we ask whether this change in sources of information translates into changes in agricultural knowledge. To do so, we examine whether AO improves farmers' ability to answer a set of 44 basic agricultural questions. The questions test the respondents on a wide range of topics, which are generally invariant to their personal circumstances.²⁵

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.

In Panel C, we find that the main ANCOVA specification does not find a statistically significant difference in measured agricultural knowledge. However, in column 2, we see that accounting for spillovers yields a modest effect (3% increase relative to the control mean at baseline) and is significant at the 10% level. Likewise, using DML increases the precision and yields a similar finding. For questions relating to the respondent's primary crop, cotton, we find that AO significantly increases knowledge by approximately 5% and is statistically significant across all specifications.²⁶

Overall, these results suggest that the AO service was successful in establishing itself as a source of information for treatment respondents in making a variety of important agricultural decisions and produced modest gains in agricultural knowledge. In the next sections, we look at whether the provision of information through AO influenced input use and agricultural productivity.

3.3. *Agricultural Input Adoption*

A number of input choices influence agricultural productivity. Cotton is the main crop grown in our sample—grown by 98.4% of the sample at baseline—and chemical inputs such as pesticides and fertilisers greatly affect cotton yields. 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. Table 4 tests whether the AO service influenced summary indices that capture recommended inputs for cotton, wheat, and cumin cultivation. The recommended inputs include seed varieties, pesticides and fertilisers (see Online Appendix Table A20 for index components).

²⁴ Online Appendix Table A3 provides crop-specific results on sources of information disaggregated by survey round. Treatment group respondents report using information from input dealers less often in making cotton pesticide decisions (−7.2% at midline), although, interestingly, they report consulting input dealers *more* often in the case of cotton fertiliser 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 fertilisers is significant at the midline.

²⁵ The full text of the questions is available in Online Appendix Table A6.

²⁶ Online Appendix Table A10 provides results on knowledge that are further disaggregated by question topic.

Table 4. *Effects on Summary Indices of Input Adoption.*

Dependent variable	Control mean (baseline) (1)	Treat–control ANCOVA (2)	Treat–control spillover (3)	Treat–control DML (4)
<i>Panel A: Input adoption by crop</i>				
Cotton management (SD units)	0.000 [1.000]	0.125* (0.065)	0.147** (0.066)	0.109* (0.064)
Wheat management (SD units)	0.000 [1.000]	0.112 (0.144)	0.090 (0.145)	0.134 (0.143)
Cumin management (SD units)	0.000 [1.000]	0.077 (0.135)	0.088 (0.136)	0.087 (0.134)
<i>Panel B: Input adoption across crops</i>				
Seed management (SD units)	0.000 [1.000]	0.091* (0.048)	0.113** (0.048)	0.077 (0.047)
Pesticide management (SD units)	0.000 [1.000]	0.050 (0.058)	0.069 (0.059)	0.043 (0.057)
Fertiliser management (SD units)	0.000 [1.000]	0.129 (0.123)	0.126 (0.124)	0.147 (0.122)
N	398	2,203	2,203	2,203

Notes: This table reports the impact of AO on agricultural knowledge, input decisions for seeds, pesticides, and fertilisers. The results use data from both the Midline survey and the Endline survey. The indices in Panel A and B aggregate information over multiple outcomes for which we expect unidirectional treatment effects. Each index consists of the average of the z-scores for each component of the index, with the control group mean and standard deviation as reference. The component scores are then weighted by the inverse of the covariance matrix of the components as in Anderson (2008). Pesticide management index: dummy to indicate purchase/use of recommended pesticides. Fertiliser management index: dummy to indicate purchase/use of recommended fertilisers. Seed management index: dummy to indicate purchase/use of recommended seeds. (See Appendix Table A20 for index details). ‘Treat’ group refers to the 802 farmers that received access to AO. Column 1 provides the mean and standard deviation for the control group at baseline. Column 2 reports the treatment effect estimate from an ANCOVA specification. Column 3 uses the ANCOVA specification and controls for the baseline treatment status of a respondent’s peer group to assess spillover effects. Column 4 reports the ANCOVA specification and uses double ML to pick an optimal set of control variables. All specifications include village fixed effects and survey round fixed effects. Asterisks denote statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.3.1. *Input adoption by crop*

Panel A of Table 4 shows that the treatment increased the cotton management index by 0.12 standard deviations in the main ANCOVA specification (see Online Appendix Table A20 for index components). This estimate rises to 0.14 standard deviations (significant at the 5% level) when accounting for spillover effects, while the DML estimate is qualitatively similar.²⁷

In contrast, we do not detect significant differences for either the indices corresponding to wheat management or cumin management across specifications.²⁸

3.3.2. *Adoption by input type*

In Panel B of Table 4 we test whether the treatment influenced input adoption as captured by summary indices that capture recommendations relating to seed, pesticide, and fertiliser management. We see that the index for seed management increases by 0.09 standard deviations (significant at the 10% level), while controlling for spillover effects increases the size of this point estimate and its precision. However, adding control covariates using DML reduces the point estimate.

²⁷ Online Appendix Table A12 suggests the reminder treatment further increased compliance with cotton recommendations.

²⁸ The standard errors also suggest that the experiment may be underpowered to detect effects for cumin (grown by just 34% of the sample). Wheat cultivation involves substantially fewer chemical inputs and is primarily done for home consumption.

In contrast, we do not find that the AO services influences the adoption of inputs relating to either pesticide or fertiliser management.

3.3.3. Discussion of input adoption results

The results above suggest that service was successful in changing the behaviour of farmers in cotton cultivation and, in particular, through influencing seed choice. In interpreting these results, we follow the framework used by Kling *et al.* (2007). In our view, since the service was intended to influence cotton farming practices, the per-comparison p -values for the effects on the overall input adoption index are appropriate. However, in Online Appendix Table A9, we adjust these asymptotic p -values for testing multiple hypotheses across a family of outcomes. In Panel B, we consider just the outcomes in Table 4. While p -value adjustments using the (conservative) Bonferroni–Holm method suggest there is a very high likelihood that the significant estimates could be observed by chance, adjusting for correlation between these outcomes using the Westfall–Young method substantially reduces the family-wise error. Nevertheless, even the Westfall–Young method suggests that one should interpret the input adoption results with some caution.

While the indices above provided the most powered tests for behaviour change, here we provide some discussion on which inputs in particular drove these results. The presence of a wide variety of cotton seeds, some counterfeit, makes seed selection a particularly important decision. In Uganda, Bold *et al.* (2017) demonstrate that low-quality inputs depress returns to hybrid seeds.

In Online Appendix Table A14 we report treatment effect estimates for components of the cotton index. We see a 6 p.p. increase (significant at the 1% level) for Vikram, a Bt cotton seed recommended by the service. Similarly, looking at pest management we find that the service increases the use of imidachlorpid by roughly 6 p.p. Imidachlorpid is a neonicotinoid pesticide that targets the nervous systems of insects and is a complement to Bt cotton which has natural defences against bollworm. We also observe that the service increases the adoption of Trichoderma, a biological method of pest control. The AO service provided extensive information in both Kharif and Rabi seasons on the use of Trichoderma as a means of preventing wilt disease in cotton and cumin.

While we do not detect increases in fertiliser adoption across all crops, looking within the index we see that treated farmers increase their purchases of ammonium sulphate by approximately 3 p.p. and NPK Grade 1 by 4 p.p. across specifications. To put these effect sizes into perspective, Duflo *et al.* (2008) found an increase of 16%–20% in fertiliser adoption in Kenya using free delivery of planting and top-dressing fertiliser, while BenYishay and Mobarak (2019) found increases of 2.2%–5.5% across treatments in pit planting and 0%–19% across treatments for composting in a study using in-person extension.

3.4. Agricultural Productivity, Input Expenditure and Profit

In Table 5, we examine agricultural productivity, demand for the AO service, and profit. Overall, we do not find evidence to suggest that AO improved cotton, wheat or cumin yields. Importantly, as rainfall is highly stochastic and an important complement to chemical inputs like fertiliser, the treatment effects may vary by season (Rosenzweig and Udry, 2020).

In Online Appendix Table A17, we see that the point estimate on cotton yields, while insignificant, is positive at midline, but negative at the endline in the ANCOVA specification (columns 2–4). On the one hand, the low temporal correlation in yields suggests that an ANCOVA

Table 5. *Effects on Yield, Demand and Profit.*

Dependent variable	Control mean (baseline) (1)	Treat-control ANCOVA (2)	Treat-control spillover (3)	Treat-control DML (4)
<i>Panel A: Yield</i>				
Cotton yield (kg/acre)	694.818 [468.751]	3.130 (14.130)	5.232 (14.314)	1.398 (13.995)
N (cotton yield)	392	2,093	2,093	2,093
Wheat yield (kg/acre)	981.132 [702.002]	-22.167 (45.237)	-29.939 (46.664)	-27.037 (43.241)
N (wheat yield)	309	681	681	681
Cumin yield (kg/acre)	172.570 [191.017]	2.142 (14.563)	5.734 (14.447)	2.141 (13.838)
N (cumin yield)	169	402	402	403
<i>Panel B: Profit and expenditure</i>				
Total input expenditure (Rupees, winsorised fraction = 0.01)	22,610.100 [18,519.387]	1,845.716** (722.518)	1,863.003** (728.563)	1,632.908** (693.432)
Profit from agriculture (Rupees, winsorised fraction = 0.01)	1.36e+05 [1.26e+05]	1915.998 (3,534.125)	2,490.148 (3,593.720)	915.223 (3,426.588)
N	398	2,203	2,203	2,203
<i>Panel C: Demand</i>				
Purchased AO subscription	0.242 [0.429]	0.060** (0.028)	0.066** (0.028)	0.060** (0.028)
N (WTP)	360	1,080	1,080	1,080
WTP in BDM game (Rupees)	71.667 [114.420]	18.188** (8.498)	21.012** (8.623)	18.151** (8.275)
N (BDM)	288	836	836	836

Notes: This table reports the impact of AO on WTP, agricultural productivity, and profit. The results use data from both the Midline survey and the Endline survey. Input expenditure includes total money spent on seeds, fertilisers, irrigation, pesticides, hired labour, and household labour priced at the mean wage of hired labour. Profit from agriculture refers to the difference between total income from all crops grown less total input expenditure in the past year. WTP reports the bid in a BDM demand elicitation game in Rupees. 'Treat' group refers to the 802 farmers that received access to AO. Column 1 provides the mean and standard deviation for the control group at baseline. Column 2 reports the treatment effect estimate from an ANCOVA specification. Column 3 uses the ANCOVA specification and controls for the baseline treatment status of a respondent's peer group to assess spillover effects. Column 4 reports the ANCOVA specification and uses double ML to pick an optimal set of control variables. All specifications include village fixed effects and survey round fixed effects. Asterisks denote statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

specification yields a more powerful test. However, baseline differences can more precisely be accounted for using a difference-in-difference specification. Indeed, doing so raises these point estimates and makes them consistently positive, though not statistically significant.²⁹

With cumin, we see a statistically significant reduction in yield at midline and an increase at endline using the ANCOVA specification. Using a difference-in-difference specification, we find a similar statistically significant increase in yield at endline, though the reduction in yield at midline is now imprecisely estimated. Similarly, we find no evidence to suggest that the treatment influenced wheat yields across seasons.

In Panel B of Table 5 we turn to how the treatment influenced input expenditure and profit. Total input expenditure includes outlays on seeds, pesticides, fertiliser, irrigation, hired labour and household labour priced at the mean wage of hired labour. We find that total input expenditure increased by nearly 8% relative to the baseline control mean across specifications. Online

²⁹ We note, however, that in Online Appendix Table A12 we find some evidence to suggest that the reminder treatment influenced cotton yields, although at the same time, decreasing wheat yields.

Appendix Table A4 breaks down the treatment effects by input type and we find that this increase is driven by a roughly 15% increase in expenditure on irrigation and expenditure on seeds which increased by roughly 20% relative to the baseline control mean. Over one-quarter of push calls contained information about weather forecasts, which provide farmers with important information on weather and reduce uncertainty about the timing and value of irrigation. In addition, chemically intensive agricultural inputs recommended by the service such as higher-yielding varieties of cotton are complements with increased irrigation.

We measure profits from agriculture by computing crop income and subtracting total input expenditure as defined above. While the point estimates for profit are consistently positive (ranging from 1.8% when accounting for spillovers to 0.6% in the DML specification, relative to profit for the control group at baseline) they are not statistically significant.

Overall, these results suggest our experiment is likely underpowered to detect effects on yield and profit, which is further complicated by measurement error. We do note here, however, that this does not suggest that the effects are very large and a more highly powered study would detect large effects. Our baseline data suggests we can detect a 11.5% increase in cotton yields and a 17.5% increase in profits with 80% power. However, as we discuss in Subsection 4.3, meta-analyses of such interventions suggest that the effects may be smaller; for smaller effects, our design has very limited power.

In the next section we discuss WTP experiments for the service and what they reveal about how respondents value the information from AO.

3.5. *Heterogeneous Treatment Effects*

In Online Appendix Table A11, we compare the combined treatment group (i.e., AOE+AO) to the control group to investigate heterogeneity with respect to respondent education and income. As such, we modify equation (2) to include a dummy for being above the median of respondent education or income and the interact of the dummy with a treatment indicator.

Treatment respondents with above-median incomes are no more likely to call into the AO line, but their total usage is nearly one hour greater than those with below-median income. Farmers with higher incomes also show differential effects in the cotton management index (about 0.24 standard deviation units higher) and wheat management index (0.56 standard deviation units higher), both of which are statistically significant at the 10% level).

In contrast, we find little evidence to suggest the treatment has a differential impact on respondents by education. The one exception is that profit is significantly higher (7% relative to the control mean) for those with below median education, but lower for those with above median education (3% relative to the control mean). As we have elsewhere suggested that our experiment is likely not powered to pick up effects on profit, we do not view these results as persuasive but rather suggestive of the need for future work.

3.6. *Spillover Effects*

Given randomisation at the household level, it is possible that access to the service indirectly influenced the outcomes of study respondents as well as those who were not a part of the study but in the networks of study respondents through information spillovers.³⁰

³⁰ In a separate paper, we document in detail how patterns of social interactions and information exchange are influenced by the AO treatment (Fernando, 2020).

Table 6. *Spillover and Peer Effects.*

Dependent variable	Control mean (baseline) (1)	Study respondents			Non-study respondents	
		Treat (2)	Fraction of Peers Treated (3)	Fraction of Peers Treated × Treat (4)	Control peer group mean (5)	Fraction of Peers Treated (6)
Called AO line	0.000 [0.000]	0.580*** (0.022)	−0.006 (0.066)	0.134* (0.080)	−	−
Total AO usage (minutes)	0.000 [0.000]	81.754*** (13.499)	12.757 (41.117)	91.001* (49.735)	−	−
Index of mobile phone-based Information usage (SD units)	0.000 [1.000]	5.461*** (0.383)	−1.133 (1.166)	1.447 (1.410)	0.000 [1.000]	−0.004 (0.069)
Proportion of cotton lost to pest attacks (%)	−	−	−	−	0.143 [0.231]	−0.036** (0.015)
Pesticide management (SD units)	0.000 [1.000]	0.056 (0.072)	−0.183 (0.220)	0.087 (0.266)	0.000 [1.000]	0.093 (0.071)
N	398	2,203			393	1,114

Notes: This table assesses whether the fraction of one’s peers assigned to the treatment group influences one’s own outcomes and how this varies for study respondents and non-study respondents. The results corresponding to study respondents (Columns 2–4) use data from both the Midline survey and the Endline survey. Column 1 reports the mean and standard deviation for the control group at baseline. Column 2 reports the coefficient on a dummy variable for receiving the AO treatment. Column 3 reports the coefficient on the number of peers who were assigned to the treatment group (fraction of peers treated). Column 4 reports the interaction between ‘fractions of peers treated’ and ‘treat’. Column 5 reports the mean and standard deviation for peers who were not respondents in the main study and who were not referenced by a treatment respondent. Column 6 reports the coefficient on the number of peers who were assigned to the treatment group, from a regression of the characteristic in question on this variable. The pest management index is as described in Panel E of Online Appendix Table A20, however the index for non-study respondents relies on data from the peer survey and usage of Imidacloprid, Acephate and Acetamiprid. All regression specification include dummies for the number of peers referenced and village fixed effects. The regressions for study respondents also include fixed effects for survey round. Asterisks denote statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6 estimates spillover effects for both study respondents and a group of ‘non-study’ respondents who were surveyed in the ‘peer survey’. As the ‘spillover’ specification in Tables 3–5 suggest, the point estimates for estimated treatment effects often increase once we account for whether peers of study respondents were also exposed to the treatment. As such, they suggest treatment respondents may have discussed advice they received with their peers. Alternatively, peers may follow suit after directly observing changes in their neighbours’ agricultural practices.

In the following section, we estimate spillover effects, using the specification in equation (4). To wit, for study respondents we amend the main ANCOVA specification (equation (2)) to include a variable that captures the fraction of peers who they listed at baseline who also received the treatment (column 3) and a control variable for the number of peers they listed at baseline. Further, we include the interaction between their treatment status and the fraction of their peers who were treated (column 4). The latter helps us test whether there are complementarities among treated respondents.

For the non-study respondents we estimate equation (4). In particular, we capture spillover effects by estimating a regression that includes a control for number of study respondents who listed this non-study respondent as a peer (i.e., the number of references) and the fraction of those references who are treatment respondents.³¹

³¹ Online Appendix Table A7 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. Controlling for baseline cotton acreage leaves the point estimates and their precision virtually unchanged.

3.6.1. *Among study respondents*

In column 4 of Table 6 we see that having a peer assigned to the treatment significantly increases the usage of AO on an extensive margin (13.4 p.p.) and on an intensive margin (91 additional minutes), suggesting that there are complementarities among respondents who have AO.

However, we find little evidence to suggest that treated peers influence information sources used in agriculture, the adoption of inputs or sowing decisions, suggesting that these outcomes are unlikely to be underestimated by a large margin.

3.6.2. *Non-study respondents*

Columns (5) and (6) of Table 6 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 for the number of peers in one's reference group. These data provide a more powerful test for spillover effects, as non-study respondents are those who are listed by a study respondent at baseline but are not themselves a part of our study. As such, non-study respondents are, by definition, exposed to study respondents.

Here, we do find some evidence for spillovers. 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. We do, however, acknowledge that this evidence is suggestive, insofar as we do not document systematic changes in pest management index among peer or, more generally, study respondent farmers. However, we find some encouragement in the fact that pest management was the most important topic covered in the service (see Online Appendix Table A1).³²

4. WTP and the Market for Information

While the results above do not allow us to unambiguously establish the effect of AO on ultimate outcomes such as yield and profits, in this section, we assess farmer WTP for the service. More generally, the financial sustainability of a subscription-based service would depend critically on users' WTP. However, markets for information may face several important challenges. Potential buyers may not be able to evaluate the quality of the goods; information may be non-rival, allowing buyers to provide the information to neighbours (which may in turn suppress demand for the service).

4.1. *Measuring WTP*

We measure WTP by offering for sale subscriptions at the conclusion of the study. We visited the 1,200 study respondents, as well as an additional 457 non-study respondents. One-quarter of this sample (chosen at random, via scratch cards) was offered a fixed price for a nine-month subscription at a single price ('take it or leave it' offer, or TIOLI); we randomly varied that price across farmers to estimate demand.³³ For the remaining three-quarters of the sample, we used a modified version of the BDM method to elicit WTP. In this method, the respondent first indicates their willingness to purchase at a series of price points. The maximum WTP is recorded, after which a randomly generated offer price is drawn. If the offer price is below the bid, the

³² Note, the peer survey did not collect information on seed choices or fertiliser use, hence we are not able to estimate these effects.

³³ The prices offered were Rs. 40 (\$0.67), Rs. 90 (\$1.5), Rs. 140 (\$2.3), Rs. 190 (\$3.2) and Rs. 240 (\$4).

sale takes place at the offer price; if the offer price is below the bid, the sale does not take place.³⁴

Overall, we find that 33% of those offered a subscription purchased it across methods (see Panel A of Online Appendix Table A16). Using the BDM method, we find that average WTP is Rs. 109 (\$1.79).³⁵ The two methods of eliciting WTP deliver similar results, with high take-up at low prices.³⁶ Approximately 85% of study respondents accepted the TIOLI price of 40, falling to 6.7% at a price of Rs. 240. Acceptance rates for these prices were similar in the BDM sample (87% at Rs. 40, 17% at Rs. 240), and 6% at Rs. 490.

Restricting attention to study respondents, in Panel C of Table 5, we find the control group is willing to pay approximately Rs. 72 (\$1.18) on average, while those in the treatment group, who have experienced the service, are willing to pay approximately Rs. 18 more (significant at 5% level). For the realised distribution of offer prices, AO users are 6 p.p. more likely to purchase a subscription. Adjusting for spillovers increases these treatment effects for both WTP (Rs. 21, significant at the 5% level) and purchasing AO (6.6 p.p. significant at the 5% level). Both of these effects suggest that exposure to the treatment resulted in an increased valuation of the service, two years after the fact.

Online Appendix Table A16 examines several correlates of demand. In addition to treatment status, we find that those who used the service more and those with more education were more likely to purchase a subscription.

4.2. Profit-Maximising Price

The distribution of WTP from the BDM exercise in our study sample enables us to calculate the revenue-maximising price a private service provider might charge. If there were approximately no cost to offering the service, a seller would maximise profits by charging Rs. 190 (\$3.12), and reach 25% of the population. At a cost of Rs. 100 (what we estimate the roughly at-scale cost to be), a provider could maximise profits at Rs. 290 (\$4.76), and reach 25% of the population. Finally, if the cost of providing the service were high (e.g., Rs. 200), profits would be maximised at a price of Rs. 490 (\$8.05).³⁷

We note that these are not true measures of the profitability of the service, as one would have to consider subscription renewal rates and customer acquisition costs, among other factors. Such an exercise is beyond the scope of this paper.

4.3. Cost-Benefit Analysis for Farmers

We estimate the costs of an AO subscription in this study at approximately Rs. 600 (\$10) per farmer per year, roughly comparable to a single in-person extension service, Rs. 517 (\$8.50) per farmer per year.³⁸ However, we note that the costs of a digital service would decline dramatically with scale with services similar to AO in India now pricing subscriptions at roughly \$2 per farmer per year. In addition, while the costs of airtime accounted for roughly 40% of the cost of provision

³⁴ The 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), and Rs. 490 (\$8.1).

³⁵ In this section and the following one, we use the conversion of 1 USD = INR 60.89.

³⁶ In Cole *et al.* (2020) we document these findings across multiple experiments in detail.

³⁷ In fact, a service provider might choose an even higher price, as we did not evaluate how demand drops when the price goes above Rs. 490 (\$8.05).

³⁸ Online Appendix Table A23 provides a breakdown of the costs of the service.

in our case, evolution in markets, such as data services in rural areas with little or no marginal cost, may considerably reduce the cost of offering services like this.

As we discuss above, our results suggest that our study is likely not powered to detect meaningful gains in yields or profits; we are therefore not able to incorporate these estimates into a cost–benefit analysis. More generally, a key challenge in evaluating low-cost interventions are the large samples necessary to identify reasonable benefit cost ratios. One approach to this power problem is to combine results from multiple studies in a meta-analysis, as done by Fabregas *et al.* (2019) in a review article. This article, surveying six studies in India and Africa that provide agricultural advice through digital technologies, finds an average impact on yields of approximately 4%, with a 95% confidence interval of (0, 0.08).³⁹

In Table 5, we estimate an increase in input costs of roughly Rs. 1,800 (\$30). As such, a 4% increase in yields would result in Rs. 4,689 (\$77) profit, net of the increase in input costs.⁴⁰ For comparison, if yields were one standard deviation higher than the mean of the Fabregas *et al.* (2019) estimate (6 p.p.) the profit would be Rs. 7,976 (\$131), while if they were one standard deviation lower, profits would be Rs. 1,461 (\$24). In each case, this exceeds the Rs. 600 (\$10) cost of the subscription.⁴¹

We note one final point: our evidence, and additional work by Fernando (2020), suggests that farmers share a significant amount of information. While this may result in reduced estimates of WTP, this also suggests the importance of studying the externalities created by such a service, as we provide some evidence for above.

5. Threats to Validity

5.1. Attrition

Online Appendix Table A8 analyses the characteristics of attritors from the study by treatment group. In the endline survey, we had 120 attritees, of which 39 were control farmers, 43 were from the AOE group, and 38 were 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. For the most part, we do not observe any significant differences between either the individual treatment arms or the combined treatment group (AOE+AO) and control group for the attritees, as measured by baseline characteristics. The one exception is in the case of wheat cultivation between the AOE group and control at midline, in which AOE attritees were less likely to grow wheat. This may in part explain why respondents in the AOE group are 8% more likely to cultivate wheat at endline relative to the control group a result we do not emphasise because it does not show a consistent pattern across treatment groups or across rounds.

5.2. Experimenter Demand Effects

A second concern is that respondents in the treatment group may offer answers that they believe the research team seeks, perhaps in the hope of prolonging the research project or due to a sense

³⁹ Further, as Fabregas *et al.* (2019) note, the marginal cost of digital agricultural interventions they survey are so low that governments might seek to invest in these programs even where estimated positive returns are noisy unless they are very risk averse.

⁴⁰ Here, we use total crop revenue for the control group at baseline: Rs.161,220 (\$2,687).

⁴¹ Were we to factor in the opportunity cost of a farmer's time spent on the AO service by using the median wage for an agricultural labourer, this would mean subtracting Rs. 76 (\$1.25) from the profit margin as the yearly use of the service amounts to roughly half a day's worth of labour.

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 supports the interpretation that our results are not driven by demand effects.

We also note that we can observe some outcomes perfectly: the AO platform records precisely how many times respondents call in. Respondents provide consistent answers to the question 'Did you call into the AO line with a question?' with a 55.5% self-reported call-in rate vs. a 53.5% call-in rate using administrative data (results not reported in tables). Finally, we note that the survey teams did not identify themselves as associated with the AO service.

5.3. Study Design

An important limitation of the study is the within-village randomisation, which introduces the possibility of information spillovers. We address this concern in two ways. First, in Tables 3–5, we include terms that account for spillover effects as in equation (3). As such, we can compare the estimated treated coefficient with (column 3) and without (column 2) spillover controls to understand the scope of the problem. For the most part, we do see some evidence that the treatment coefficients are lower when we don't account for spillover effects, though these differences are not typically economically large.

Second, we estimate spillover effects directly by comparing control group respondents who are exposed to peers who are in the treatment group, to those who are not, as in Table 6. In Online Appendix Table A24, using the same specification as in Table 6 (columns 2–4), we find that for all input indices, the coefficient on 'fraction of peers treated' (column 3), which captures the effect of exposure to a treated peer for the control group, is rarely different from zero. We believe future research, using a clustered experimental design that includes 'pure control' villages, could help more definitively address the importance of spillovers.

6. Conclusion

This paper presents the results from a randomised 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.

First, we show that the intervention was successful in generating a substantial amount of AO usage, with roughly 60% of all treatment respondents calling in to listen to content or ask a question within seven months of beginning the intervention, and nearly 90% using it after two years. We then showed that AO had a large impact on reported sources of information used in agricultural decisions.

Having established AO as a reliable source of information, we then show that advice provided through AO resulted in farmers consistently changing their input decisions in cotton cultivation and, in particular, their seed choices in line with recommendations from the service.

We do not find evidence to suggest the service improved yields or profits. Unfortunately, the imprecision of our estimates suggests the study was underpowered to detect effects for these outcomes, suggesting a need for future research. In particular, the study of on-demand interventions will need to carefully consider the challenge in detecting treatment effects presented by a non-homogeneous treatment that is a function of important agricultural shocks like pest infestations and weather.

These results suggest the importance of further research on the importance of information in smallholder agriculture. Many important questions remain unanswered, such as: the importance of the customised on-demand vs. 'push' information; the role of existing trust in the NGO associated with the service in promoting credibility; and the importance of the 'unbiased' nature of the service relative to commercially motivated information provided. We hope to test these in future experiments.

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.

Harvard Business School, USA
University of Notre Dame, USA

Additional Supporting Information may be found in the online version of this article:

Online Appendix **Replication Package**

References

- Aker, J. (2011). 'Dial "a" for agriculture: using ICTs for agricultural extension in development countries', *Agricultural Economics*, vol. 42(6), pp. 631–47.
- 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*, vol. 3, pp. 2343–78.
- Anderson, M.L. (2008). 'Multiple inference and gender differences in the effects of early intervention: a reevaluation of the Abecedarian, Perry Preschool, and early training projects', *Journal of the American Statistical Association*, vol. 103(484), pp. 1481–95.
- 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*, vol. 3(4), pp. 186–214.
- Belloni, A., Chernozhukov, V. and Hansen, C. (2014). 'Inference on treatment effects after selection among high-dimensional controls', *Review of Economic Studies*, vol. 81(2), pp. 608–50.
- BenYishay, A. and Mobarak, A.M. (2019). 'Social learning and incentives for experimentation and communication', *Review of Economic Studies*, vol. 86(3), pp. 976–1009.
- Birkhaeuser, D., Evenson, R. and Feder, G. (1991). 'The economic impact of agricultural extension: a review', *Economic Development and Cultural Change*, vol. 39(3), pp. 607–50.
- Bold, T., Kaizzi, K., Svensson, J. and Yanagizawa-Drott, D. (2017). 'Lemon technologies and adoption: measurement, theory and evidence from agricultural markets in Uganda', *Quarterly Journal of Economics*, vol. 132(3), pp. 1055–100.
- Casaburi, L., Kremer, M., Mullainathan, S. and Ramrattan, R. (2020). 'Harnessing ICT to increase agricultural production: evidence from Kenya', Mimeo.
- Cole, S., Fernando, A.N., Stein, D. and Tobacman, J. (2020). 'Field comparisons of incentive-compatible preference elicitation techniques', *Journal of Economic Behavior and Organization*, vol. 172, pp. 33–56.
- Cole, S., Giné, X., Tobacman, J., Topalova, P., Townsend, R. and Vickery, J. (2013). 'Barriers to household risk management: evidence from India', *American Economic Journal: Applied Economics*, vol. 5(1), pp. 104–35.
- Cole, S. and Sharma, G. (2020). 'The promise and challenges of implementing ICT in Indian agriculture', Mimeo.
- Cole, S.A. and Fernando, A.N. (2012). 'The value of advice: evidence from mobile phone-based agricultural extension', Mimeo.
- Draca, M., Sadun, R. and Van Reenen, J. (2007). 'Chapter 5: Productivity and ICTs: A Review of the Evidence', *The Oxford Handbook of Information and Communication Technologies*, pp. 100–47.

- Duflo, E., Kremer, M. and Robinson, J. (2008). 'How high are rates of return to fertilizer? Evidence from field experiments in Kenya', *American Economic Review*, vol. 98(2), pp. 482–88.
- Duflo, E., Kremer, M. and Robinson, J. (2011). 'Nudging farmers to use fertilizer: theory and experimental evidence from Kenya', *American Economic Review*, vol. 101(6), pp. 2350–90.
- Fabregas, R., Kremer, M., Lowes, M., On, R. and Zane, G. (2020). 'SMS-extension and farmer behavior: lessons from six RCTs in East Africa', *Mimeo*.
- Fabregas, R., Kremer, M. and Schilbach, F. (2019). 'Realizing the potential of digital development: the case of agricultural advice', *Science*, vol. 366(6471), p. eaay3038.
- Fafchamps, M. and Minten, B. (2012). 'Impact of SMS-based agricultural information on Indian farmers', *World Bank Economic Review*, vol. 26(3), pp. 383–414.
- Feder, G. (2005). 'The challenges facing agricultural extension—and a new opportunity', *New Agriculturalist*, <http://www.new-ag.info/en/view/point.php?a=1253>.
- Feder, G., Lau, L. and Slade, R. (1987). 'Does agricultural extension pay? The training and visit system in Northwest India', *American Journal of Agricultural Economics*, vol. 69(3), pp. 677–86.
- Fernando, A.N. (2020). 'Seeking the treated: the impact of mobile extension on information exchange in farmer peer groups', *Mimeo*.
- Foster, A. and Rosenzweig, M. (2010). 'Microeconomics of technology adoption', *Annual Reviews of Economics*, vol. 2(1), pp. 395–424.
- Gandhi, R., Veeraraghavan, R., Toyama, K. and Ramprasad, V. (2009). 'Digital green: participatory video and instruction for agricultural extension', *Information Technologies and International Development*, vol. 5(1), pp. 1–15.
- Giné, X. and Yang, D. (2009). 'Insurance, credit, and technology adoption: field experimental evidence from Malawi', *Journal of Development Economics*, vol. 89(1), pp. 1–11.
- Glendenning, C., Babu, S. and Asenso-Okyere, K. (2010). 'Review of agricultural extension in India: are farmers' information needs being met?', IFPRI Discussion Paper 01048.
- Jack, K. (2011). 'Market inefficiencies and the adoption of agricultural technologies in developing countries', White paper prepared for the Agricultural Technology Adoption Initiative, J-PAL (MIT)/CEGA (Berkeley).
- Kling, J.R., Liebman, J.B. and Katz, L.F. (2007). 'Experimental analysis of neighborhood effects', *Econometrica*, vol. 75(1), pp. 83–119.
- Lobell, D.B., Azzari, G., Marshall, B., Gourlay, S., Jin, Z., Kilic, T. and Murray, S. (2020). 'Eyes in the sky, boots on the ground', *American Journal of Agricultural Economics*, vol. 102(1), pp. 202–19.
- McKenzie, D. (2012). 'Beyond baseline and follow-up: the case for more t in experiments', *Journal of Development Economics*, vol. 99(2), pp. 210–21.
- Muralidharan, K., Romero, M. and Wüthrich, K. (2020). 'Factorial designs, model selection, and (incorrect) inference in randomized experiments', NBER Working Paper 26562.
- Patel, N., Chittamuru, D., Jain, A., Dave, P. and Parikh, T. (2010). 'Avaaj 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*.
- Rosenzweig, M.R. and Udry, C. (2020). 'External validity in a stochastic world: evidence from low-income countries', *Review of Economic Studies*, vol. 87(1), pp. 343–81.
- Saito, K.A. and Weidemann, C.J. (1990). *Agricultural Extension for Women Farmers in Africa*, 103, World Bank Publications.
- Telecom Regulatory Authority of India. (2015). 'Annual report 2014–2015 telecom regulatory authority of India'.
- Van Campenhout, B. (2017). 'There is an app for that? The impact of community knowledge workers in Uganda', *Information, Communication & Society*, vol. 20(4), pp. 530–50.
- World Bank Group. (2012). *World Development Indicators 2014*, World Bank Publications.