One Size Fits All?
Experimental Evidence on the Digital Delivery of Personalized Extension Advice in Nigeria*

Aminou Arouna[a], Jeffrey D. Michler[b], Wilfried G. Yergo[a], & Kazuki Saito[a]

[a] *Africa Rice Centre (AfricaRice), Bouake, Côte d'Ivoire*
[b] *Department of Agricultural and Resource Economics, University of Arizona, Tucson, USA*

September 2020

Abstract

Blanket advice on optimal fertilizer application rates has failed to achieve potential yield gains for crop production in much of Sub-Saharan Africa. However, digital technology now makes it possible to deliver personalized extension services to farmers at a much lower cost. We present results from a randomized control trial designed to evaluate the effectiveness of a mobile application (or app) that provides personalized advice on rice nutrient management. We find that households who were only given personalized advice increase their yield by seven percent and increase their profit by 10 percent. We show that, on average, personalized advice increases yields without increasing the overall quantity of fertilizer used. We conclude that the scaling of personalized extension services could improve productivity and livelihoods in Sub-Saharan Africa without necessarily increasing the total amount of fertilizer in use.

JEL codes: C93, D24, O33, Q16
Keywords: information interventions, extension, information and communication technology, decision support tools, RiceAdvice

*Corresponding author email: jdmichler@arizona.edu. The authors would like to thank the Global Rice Science Partnership (GRiSP), the CGIAR Research Program on Rice Agri-Food Systems (RICE), the African Development Bank (AfDB) project "Support to Agricultural Research for Development of Strategic Crops in Africa" (#2100155022217), and the Government of Japan for providing financial support for the experiment and data collection.

All environments are intractably local.

-James C. Scott, *Seeing Like a State*

Throughout Sub-Saharan Africa, efforts by governments and development agencies to spur agricultural intensification have been meet with continued low levels of adoption of improved inputs. One explanation for the lack of adoption is that returns to improved inputs are highly heterogeneous. Unable to represent the complexity of actual farms, institutions produce official, standardized recommendations for the levels of input use, and the resulting expected yields, that are radically simplified. These simplified recommendations, as Scott (1998) notes, are made for administrative convenience, not ecological considerations. The result are development programs and policy "solutions" that pay scant attention to the particularity of a given field or farmer, collapsing or ignoring distinctions that might otherwise be relevant. While the recommendation to adopt improved inputs is accurate for the stylized farm, the necessary input levels, and the subsequent returns, can be very different on any particular farm.

To some extent, the simple abstractions made by institutions in promoting the adoption of improve inputs are understandable. Historically, it was prohibitively expensive to account for heterogeneity in fields and farms when producing recommendations on optimal input levels. This made blanket recommendations necessary, even if failure to formulate advice that was soil, crop, and climate specific resulted in inefficiencies, reducing yield and profit, and generated negative environmental externalities. In the past few years, however, mobile technology in the form of decision support tools (DSTs) have greatly reduced the cost of delivering personalized extension services (MacCarthy et al., 2017; Tjernström et al., 2020). DSTs allow for farmer input regarding local conditions in both the household and the environment. The goal of DSTs is to reduce inefficiencies coming from grossly simplified recommendations, thereby raising the productivity and profitability from adopting improved inputs, though there is little rigorous evidence of their effectiveness in this regard.

We assess the impact of personalized extension services delivered using a specific DST: RiceAdvice. RiceAdvice is an Android-based application (or app) that was developed by AfricaRice to provide personalized recommendations on nutrient management (type, quantity, and timing of fertilizer) in rice production. The app utilizes information and communication technology (ICT) that enables extension agents to gather data from farmers about the farmers' particular local context. The app then provides farmers with specific crop, field, and seasonal advice regarding fertilizer application and agro-management practices (Saito et al., 2015a).

To measure how households respond to personalized advice from the app compared to blanket advice from extension agents, we conduct a clustered randomized control trial (RCT) in Nigeria. We establish two simple treatment arms: 1) rice production with personalized advice on nutrient management and 2) rice production with personalized advice plus a grant that provides the recommended level of fertilizer. The second treatment aims to assess the importance of liquidity constraints on adoption of improved inputs. In addition to the two treatments, a control group received the government's standardized advice regarding fertilizer application rates. We calculate impacts using OLS along with Analysis of Covariance (ANCOVA) estimation. Regardless of the estimation strategy, we find that households who were only given personalized advice increase their yield by around seven percent and increase their profit from rice by around 10 percent. Households who received personalized advice in combination with the provision of fertilizer increase their yield by around 20 percent and increase their profit by around 23 percent. Interestingly, we find that RiceAdvice tends to decrease the average amount of fertilizer used by those in the treatment. In response to the recommendations provided by the app, households decrease their use of NPK fertilizer while keeping their use of urea fertilizer unchanged. This suggests that personalized recommendations, via DSTs, could improve productivity and livelihoods in Sub-Saharan Africa relative to existing blanket recommendations, without necessarily increasing the overall amount of chemical fertilizer use, and the corresponding negative effects on the environment.

Our study contributes to the extensive literature on technology adoption in several ways. First, we provide evidence that failure to account for heterogeneity in soil quality may be a limiting factor in the adoption of improved technology, at least among rice farming households in Nigeria. Numerous studies have documented that households in developing countries fail to adopt apparently profitable technologies (Foster and Rosenzweig, 2010; Jack, 2011, Magruder, 2018). Reasons for the lack of adoption include credit and risk constraints (Karlan et al., 2014), insecure property rights (Burchardi et al., 2019), limited access to insurance (Casaburi and Willis, 2018), missing input markets (Byerlee and Deininger, 2013), missing output markets (Michler et al., 2019), learning externalities (Conley and Udry, 2010), and risk preferences (Liu, 2013). Suri (2011) has suggested that unobserved heterogeneity may make the returns to adoption for any individual farmer unprofitable, even if average returns to the technology are positive. While Suri (2011) focuses on heterogeneity in farmer ability, we focus on heterogeneity in soil quality. We find that providing personalized advice on optimal input use that accounts for differences in soil quality results in higher yield and profit, when compared to blanket extension advice. This suggests that how individuals are taught to use the technology plays an important role in whether that technology is profitable or not.

Second, we present new evidence regarding the effectiveness of information interventions. In general, information-only interventions have yielded null results. Bettinger et al. (2012) study the impact of providing aid eligibility information to low-income households, but find no effect. Ashraf et al. (2013) find that providing information about water purification fails to increase demand for pure water. Bryan et al. (2014) find no effect of information about migrant opportunities on the decision to migrate. One reason why studies of information-only interventions frequently find null results may be due to the overly broad nature of much of the information tested in these interventions. In the information-only arm of the RCT, we find positive and significant effects of personalized extension services on yield and profit. A study similar to ours, Tjernström et al. (2020), tests a mobile game designed to provide personalized information on input use for maize farmers in Kenya and finds similarly positive results.

Third, our focus on delivering extension services in Nigeria adds to the small but growing literature on ways to improve the effectiveness of agricultural extension in developing countries. Previous research has looked at the effect of increased supervision on production in Madagascar (Bellemare, 2010), gender and network effects on the adoption of conservation agriculture in Mozambique (Kondylis et al., 2016; Kondylis et al., 2017), the effect of extension combined with grants on livestock production in Uruguay (Mullally and Maffioli, 2016), the impact of extension on food security in Uganda (Pan et al., 2018), and the use of ICTs to deliver blanket advice to farmers in India, Kenya, and Uganda (Cole and Fernando, 2018; Casaburi et al., 2019; Van Campenhout et al., 2020). To date, this literature has focused solely on improvements to the delivery of generic advice. One exception is Krishnan and Patnam (2014), who contrast learning from extension agents with learning from neighbors in Ethiopia. Though not its focus, Krishnan and Patnam (2014) provide indirect evidence for Scott's (1998) argument that personalized, localized advice is far more effective than centrally planned, broadly accurate advice. Our paper provides direct evidence in support of Scott's (1998) argument.

Finally, we present some of the first experimental results on the use of DSTs and other mobile technologies to address barriers to adoption. Innovations in ICT have greatly reduced the cost of delivering information that is targeted to individual users. While personalization is increasingly used to provide advertising content to internet users, the innovations are rapidly being applied to a host of new purposes, including education (Walz and Detering, 2015). Recent studies have examined the use of DSTs in adapting to climate change (Watkiss et al., 2015), operationalizing ecosystem services (Grêt-Regamey et al., 2017), and improving agricultural production (Rose et al., 2016). However, most studies of DSTs are based on either observational data or on data collected from highly controlled laboratory-type settings. To our knowledge, this study, along with Tjernström et al. (2020), are the first RCTs to assess the economic impact of DSTs on agricultural production.

In what follows, we first detail the setting and technology (Section 1). We then describe the experimental design and data collection process (Section 2). Section 3 outlines the assumptions underlining the causal pathway and our resulting empirical strategy. We present primary and secondary results in Sections 4 and 5. Section 6 discuss implications of the study and concludes.

1. Rice in Nigeria and the RiceAdvice App

Rice now represents the staple food for more than 750 million people in Sub-Saharan Africa (USDA, 2018). Nigeria, a country with 170 million people, has a population growth rate of 2.5 percent per year, while rice consumption has risen at approximately six percent per year. This makes Nigeria the top consumer of rice in Sub-Saharan Africa.

Rice production in Nigeria is concentrated in seven states in the northwest of the country (Kano, Kaduna, Jigawa, Sokoto, Zamfara, Kebbi, and Niger) where 72 percent of rice is produced. Although rice production is increasing, local production represents only 55 percent of consumption (Saito et al., 2015b). As a result, Nigeria imported nearly 2.6 million tons of milled rice at a cost over one billion U.S. dollars in 2017 (USDA, 2018). The gap between production and consumption is partly due to yields that are well below their potential. Average yield is around two tons per hectare, while the potential yield for water-unlimited lowland rice is up to 12 tons per hectare (van Oort et al., 2017). With rice yield gaps that range from 10 to 70 percent in Sub-Saharan Africa, Nigeria is among the countries with the largest difference between potential and actual yields (Saito et al., 2015b).

To reduce its reliance on imported rice, the government of Nigeria has embarked on a program to increase production and productivity through intensifying rice cultivation. Among other actions, the government recently launched the Growth Enhancement Support Program, a major policy shift that transfers the supply system for farm inputs from the state to the private sector. However, the policy aims at increasing adoption of fertilizer by addressing only missing input or output markets. The effectiveness of such a policy may be limited unless it also seeks to address on-farm inefficiency in fertilizer use due to heterogeneity in soil quality. To help address this gap, AfricaRice, in conjunction with national partners, developed the RiceAdvice mobile app.

The RiceAdvice app is an Android-based DST that extension agents can use to provide farming households with pre-season, field-specific management guidelines for rice production. The extension advice includes a nutrient management plan, a suggested crop calendar, and information regarding best practices for rice cultivation. To generate this advice, farmers provide information on the geographic location of the plot, descriptive soil quality measures, local rice-growing conditions, seed variety, typical management practices, expected sowing date, availability of fertilizers, market prices for inputs, and expected production costs (Saito et al., 2015a). Figure 1.A provides examples of the data input screens for the app. As output, RiceAdvice gives field-specific information on the chemical fertilizers required, a fertilizer application plan, fertilizer cost, and recommendations regarding cultivation practices, such as levelling, timely and uniform sowing, weeding, and anticipated harvest date. Figure 1.B provides examples of the personalized output from the app. Saito et al. (2015a) provides more detail on the specifics of the app.

2. Experimental design, sampling, and data

To assess the impact of personalized advice provided by the DST on household decision making, we conducted a clustered RCT in the main rice producing state of Kano. Households were randomly assigned to a control group or one of two treatment groups. An extension agent then visited each farmer (both control and treatment) one time. The same set of extension agents were used in both treatment and control to diminish agent effects. The visit was done physically and it occurred at the beginning of the growing season. In the control group (C), households received blanket advice provided by the extension agent.[1] The blanket advice, which comes from the Federal Ministry of Agriculture and Rural Development, varies solely by crop and whether soil is classified by low, medium, or high fertility (Chude et al., 2012). Table A.1 in the Appendix reproduces the government's recommendations for fertilizing rice. In the first treatment group (T1), farming households were offered personalized advice delivered by the extension agent using the RiceAdvice app. This information-only treatment was designed to assess whether the

[1] Using a pure control group that received no contact from extension would confound the effects of the extension agent with the effects of the app. One would be unable to tell if the app itself was having an effect or if it was just the effect of a visit from an extension agent. Thus, the experiment was designed to allow identification of the marginal improvement of the app relative to the type of extension currently provided in Nigeria.

personalized advice was more effective than the blanket advice offered by extension agents. In the second treatment group (T2), farming households were offered personalized advice using RiceAdvice along with a 100 percent subsidy (grant) for the quantity of fertilizer recommended by the RiceAdvice app. The grant was provided in-kind, with AfricaRice delivering the recommended amount of each type of fertilizer to each household. This subsidy treatment was designed to remove the liquidity constraint that is commonly assumed to be binding for smallholder farmers in developing countries.[2]

2.1.Sampling, compliance, and attrition

To select the study area and farming households in the sample, we used a multi-level stratified sampling approach. First, we selected Kano state because it is the major rice producing region in Nigeria. Second, we identified the rice producing Local Government Areas (LGAs) in Kano, and randomly selected five from the eight major irrigated rice production LGAs (see Figure A.1 in the Appendix). Third, we stratified by LGA so that the random selection of rice producing villages would be proportional to the total number of rice-growing villages in the LGA. In total, 35 villages were selected and were randomly divided into two groups: 18 treated villages and 17 control (see Figure 2). In addition, the treated villages were randomly assigned to one of the two treatment groups: treatment villages that received the information-only treatment (T1) and treatment villages that received the subsidy treatment combined with RiceAdvice (T2). Finally, within villages we randomly selected 20 households from a census of all rice farming households. In total, 700 households were sampled in 35 villages, including 360 treated households and 340 control

[2] Note that we do not consider the two treatment arms as part of a partial 2x2 design but rather as independent treatments, which we evaluate as such.

households. The treated households were divided into two groups: 260 treated households for T1 and 100 treated households for T2.

The sample size of each group was determined by our own power calculations and the administrative budget available to provide the fertilizer subsidy for households in the second treatment group.[3] To increase the power of the sampling, we selected treatment and control villages in each LGA using a matched pair randomization approach (Imai et al., 2009). However, implementation of our design was imperfect in that one control village was treated by an extension agent, which caused 20 households assigned to the control group to be given the information-only treatment. This represents a contamination rate of 2.8 percent but analysis shows that this has no effect on the significance of RiceAdvice's impacts on our outcomes.[4] The rate of non-compliance was also low. Only 12 households out of 700 did not use the personalized advice, which represents a rate of 1.7 percent. This non-compliance rate is much lower than those experienced in other information interventions. For instance, Fafchamps and Minten (2012) report a non-compliance rate of 27 percent in their SMS-based information intervention. The main reason given for non-compliance in our sample was uncertainty about the riskiness of applying fertilizer at a rate different from a household's historic application rate. Despite the low non-compliance rate, we still interpret our results as intent to treat (ITT) effects.

We use three rounds of a household-level panel survey data in our analysis. First, a baseline survey was conducted in early 2016 in order to collect information on farm production before the treatment. We then conducted our intervention ahead of the rice growing season. A second survey

[3] To calculate the sample size, we used rice production data from 200 households who participated in an on-farm trial in the survey area. With a minimum detectable effect size of 0.5t/ha (the yield control was estimated to be 4t/ha and 4.5t/ha for the treatment groups), a standard deviation of 1.64 t/ha and a power of 0.8, we required sample sizes of 340 (C) and 340 (T1+T2) to detect effects at standard levels of confidence.

[4] The impact on outcomes do not change substantially when the contaminated households are dropped from the analysis or included in the treatment. The results are not shown here but can be obtained from the authors.

was conducted immediately after the 2016 rice harvest was completed. Finally, we conducted a follow-up survey one year later, at the end of the 2017 rice season in order to analyze the behavior of households during the second year following the intervention. While we were able to follow-up with all households in every year, not every household chose to produce rice in every year. As an example, six of the 700 household chose not to produce rice in 2016, which is less than one percent of the sample. Thus, while there is no classical attrition in the data, sample size does vary slightly across years. Table A.2 in the Appendix reports these differences in sample size. We also follow Lee (2009) in calculating bounds for how large an effect these missing values might have on our estimates.

2.2.Measurement

In order to reduce both noise and bias in the measurement of our outcome variables, we relied as much as possible on objective instead of self-reported information.[5] To determine yields, rice plots were traced using hand-held GPS devices. At harvest, we implemented one-meter squared crop cutting and took crop cuts from two locations in each plot. The quantity of fertilizer input was measured on a scale for those receiving the fertilizer subsidy but was self-reported for those in the information-only treatment and control households. Rice income is calculated by multiplying household yield (in tons per hectare) by the average unit price of paddy rice (in US$ per ton) in the data. Rice profit is simply the difference between rice income and the sum of all rice production costs, including the cost of the subsidy, while excluding the cost of labor and equipment, for which unit prices are notoriously difficult to calculate.

[5] Recent research has shown that self-reported measurements of area planted and quantity harvested can be susceptible to non-classical measurement error (Abay et al., 2019). Failing to account for this mis-measurement can bias results (Gourlay et al., 2019).

Socio-economic data was collected through household interviews. Measurement of age and household size are straightforward. Education is recorded as a binary variable equal to one if the farmer had received formal education for at least six years (completed primary school). The household's main activity is measured by a binary variable equal to one if crop production is the main occupation of the household head. In cases where farming is not the main activity, the household head is typically engaged in trade or transportation, with other household members responsible for the farm. The number of agricultural training days is measured as the number of days the farmer participated in agricultural training over the previous twelve-month period. Access to credit is a binary variable equal to one if the household received credit to cover the cost of any farm production practice, not just rice, over the last twelve months.

2.3.Baseline balance checks

Table 1 presents the pre-treatment balance of baseline characteristics based on the original, pre-contaminated randomization.[6] Column (1) reports the mean value of each variable for the control group and its standard deviation, while columns (2) - (4) report the coefficients from OLS regressions comparing treated households with the control. We regress the variable of interest (row) on an indicator of treatment status (column) along with LGA (strata) fixed effects and standard errors clustered at the village level. Column (2) shows the pre-treatment difference in the means between assignment to any treatment (T) and the control group (C). Column (3) shows the difference between the information-only treatment group (T1) and the control group. The pre-treatment difference between the subsidy group (T2) and the control group is in column (4).

[6] We also check for balance based on the post-contaminated, random assignment (see Table A.3 in the Appendix). As is to be expected if the contamination occurred by chance, the balance between treatment and control in the two samples (pre- and post-contamination) is virtually identical.

The coefficients suggest a good balance for most household characteristics. The two exceptions are age of household head, where those in the control are older by two years compared to those in the information-only treatment, and the area planted to rice, where those in either of the treatment groups cultivate a quarter of a hectare more than those in the control group. Among the variables related to farm production, there is balance on all the outcome variables except when comparing the subsidy treatment (T2) to the control. At baseline, households in the second treatment group applied 25 kg/ha more urea than households in the control group. We account for these differences in baseline characteristics in several of our econometric specifications.

3. Empirical Framework

3.1. Causal pathway and outcomes

Before outlining our empirical strategy, it is useful to formulate the assumptions underlying the causal pathway from the provision of personalized extension advice via DST to measured outcomes. First, we assume that households are aware that nutrient management (type of fertilizer, quantity of fertilizer, and application timing) has a direct effect on rice yield and profitability. Second, we assume that households believe that personalized advice from DSTs such as RiceAdvice will recommend levels of input use that differ from their current level of input use. Third, we assume that treated households believe that the use of personalized advice will increase the productivity of rice compared to blanket advice. Finally, we assume that the treated and non-treated households sell rice at the same price and that the value of the increase in yield will be greater than the change in costs related to the use of the personalized advice. These assumptions are supported by anecdotal evidence in Saito et al. (2015a) and MacCarthy et al. (2017).

Based on these assumptions, the causal pathway is straightforward. RiceAdvice will generate personalized advice that differs from the household's current practice and thus will result in a change in the quantity, type, and timing of fertilizer. This change in nutrient management will affect land productivity, leading to changes in production. Ultimately, the changes in production will have a positive impact on profit.

Given the causal pathway, our main outcomes of interest are rice yield and rice profit. However, we also report on secondary outcomes in order to elucidate the causal chain. In the results, we explicitly distinguish between the expected main outcomes and secondary outcomes (the quantity and type of fertilizer and application timing).

3.2.Intent-to-treat (ITT) estimation

We focus on the estimation of the impact of personalized advice regarding nutrient management on rice yield and rice profit. To estimate these impacts, we compare the outcomes for households in each treatment group with the outcomes in the absence of the treatment. Because we have the benefit of observing each household in our sample before and after treatment, we employ two different methods to calculate the intent-to-treat (ITT) effects, which measure the effect of living in a village randomly assigned to a treatment, irrespective of actual household participation in the treatment. These methods are i) an OLS estimate that uses only the post-intervention data and ii) an Analysis of Covariance (ANCOVA) estimate that uses the baseline and post-intervention data. We estimate the ITT effect for both methods with and without baseline control variables.

For the post-intervention data, we use OLS to estimate the ITT effect (ρ_i) for household h in village v and LGA g as:

$$S_{hvg} = \alpha + \rho_1 T1_h + \rho_2 T2_h + \sigma_g + \varepsilon_{hvg} \tag{1}$$

where S_{hvg} is the observed outcome variable, $T1$ and $T2$ are household-level indicators that equal one depending on which treatment group the household was assigned to and is zero otherwise. Additionally, σ_g is a strata fixed effect that accounts for variation across the LGAs and ε_{hvg} is an idiosyncratic error term that is orthogonal to the ITT effect because of the randomization. In order to account for the imbalance in some of the baseline characteristics, we specify a second OLS regression that adds a vector of household covariates. These covariates include age of the household head, household size, the number of days the farmer participated in agricultural training, and indicators if the household head has any formal education, if crop production is the main household occupation, and if the household had access to credit over the last twelve months.

Our second estimator is an ANCOVA estimate of the treatment effect:

$$S_{hvg,t} = \alpha + \gamma_1 T1_{h,t} + \gamma_2 T2_{h,t} + \mu S_{hvg,PRE} + \delta_t + \sigma_g + \varepsilon_{hvg,t}. \quad (2)$$

Here $S_{hvg,PRE}$ is the value of the outcome variable from the pre-treatment growing season and γ_i is the coefficient on the ANCOVA estimate of the ITT effect for each treatment group. The equation also includes time fixed effects (δ_t) in addition to strata fixed effects. The ANCOVA estimator has more power than the typical difference-in-difference estimator, especially when there are multiple rounds of post-treatment data (McKenzie, 2012), which we have in our sample. Similar to the OLS estimates, we also estimate ANCOVA with and without covariates.

3.3.Sampling weights, clustering, and multiple hypothesis testing

Because we used a multi-level stratified sampling approach, different households have different probabilities of being sampled. As a result, assuming equal probability could lead to biased estimates of the population effects (Ksoll et al., 2016). Therefore, we use sampling weights calculated as the inverse probability of being selected in any given village for each observation. We use the weighted data in all the regressions, though our results are robust to using the raw data.

Because the ultimate sampling units (households) are clustered within our unit of randomization (village), we cannot rule out serial correlation within a village. Although the intracluster correlation coefficient (ICC) is relatively low (see Table A.4 in the Appendix), ignoring the clustered design will lead to standard errors that are too small and t-statistics that are too large. Even when individual behaviour may generate homoskedastic regression functions within a cluster, there is heterogeneity between villages, and there will be heteroskedasticity in the overall regression (Cameron and Miller, 2015). Therefore, we use heteroskedasticity robust-standard errors clustered at the village level for all inference.

Because we are making inference on a large number of hypotheses, it is possible that significant results emerge from our analysis due to chance rather than actual treatment effects. While the problem of multiple inference is well known, there is as yet no consensus regarding the best way to correct for multiple hypothesis testing. We follow Arouna et al. (2019) and adjust the p-values in a number of different ways. We calculate Romano-Wolf adjusted p-values following Clarke et al. (2019) to correct for the familywise error rate (FWER), the probability of making at least one false discovery among a family of comparisons. We also calculate sharpened q-values as in Anderson (2008) to correct for the false discovery rate (FDR), the probability of making at least one false discovery among the discoveries already made. In the Appendix, Table A.5 and Table A.6 present the results of these corrections along with the unadjusted p-values from standard errors clustered at the village level. Our findings are generally robust to the correction for multiple hypothesis testing and we highlight were differences exist.

4. Results for primary outcomes

Before proceeding to the regression results, we first present graphical evidence of the simple mean difference between each treatment and the control. Figure 3 draws the distribution of post-treatment outcomes pooling both 2016 and 2017 data. To these kernel density plots we add vertical lines to mark the unconditional mean for each group for each outcome. Visual inspection reveals substantially larger means for each treatment group relative to the control group for yield and profit. The graphs also reveal that the means for the information-only and the subsidy treatment are fairly similar. In terms of fertilizer use, households in the treatment groups tend to use less fertilizer than those in the control group, though these differences are not as pronounced as those in yield and profit. This suggests that RiceAdvice may be having an effect on not only the rate of fertilizer use but the timing as well. In the remainder of this section we present the results of the primary outcomes (yield and profit) and in the next section examine whether the intervention also changed input management practices.

For each outcome we begin by presenting OLS and ANCOVA estimates of the ITT for the pooled post-treatment data. The odd numbered columns are without covariates while the even numbered columns include covariates. The ANCOVA regressions also include the pretreatment outcome (2015) as a covariate. Because they rely on the full sample of data, and control for baseline outcomes, the ANCOVA estimates with covariates are our preferred estimates. We then estimate effects for each year individually, using OLS and ANCOVA, in order to gain a better sense of how outcomes change over time.

4.1. Treatment effect on yield

Table 2 presents the ITT effects of personalized advice from the RiceAdvice DST on rice yield by including indicator variables for each treatment group (information-only and information with the subsidy). We then compute point estimates and standard errors for the linear combination of the treatment coefficients to determine the combined effect of the treatments. We also calculate estimates of the difference between the two treatment groups to determine the effect of adding the subsidy to the information-only treatment.[7]

We find consistent evidence that personalized advice increases rice yield subsequent to the treatment (Panel A). For those in the information-only treatment, yields increase by about 250 kg/ha, which represents an increase of seven percent compared to the control. For households that received the fertilizer subsidy in addition to RiceAdvice, yields increase by about 730 kg/ha, which represents a 20 percent gain over yields for control households. Not only is the effect size of T2 larger than T1, when we compare outcomes between these two treatment arms, we find that this difference is statistically significant. The similarity between OLS and ANCOVA estimates, as well as between regressions with and without controls, suggests that the small differences in baseline characteristics are uncorrelated with treatment.

While we give preference to the ANCOVA results using all three years of data, it is useful to understand how outcomes change over time. To do this, we present results in Panel B that rely only on outcomes in 2016, the harvest immediately following the intervention. Across all regressions we find effect sizes in 2016 that are similar to the estimated effect sizes using all rounds of data. Panel C presents estimates of the ITT using yield data from one year after the intervention (2017). Given that soil characteristics change very slowly over time, the advice provided as part of the study should still be valid a year later. Thus, we would expect impacts to be of a similar size regardless of whether we use the 2016 or 2017 yield data. Treatment effects for the group that received the subsidy are larger than the overall effect but again, not substantially so. The treatment effects for the information-only group in 2017 are not significant despite point estimates that are

[7] It should be noted that since our experiment is not a 2x2 design, we are unable to estimate the overall effect of RiceAdvice. We can only estimate the effect of just RiceAdvice by itself and the effect when RiceAdvice is combined with the subsidy.

similar to those for 2016 and overall. We believe the lack of significance is due to the increase in the standard errors on the coefficients and not due to a reduction in the impact of the treatment over time.

These results imply that knowledge matters, even without relaxing the liquidity constraint. If liquidity were the primary constraint limiting agricultural intensification, then the information-only treatment would have no discernable effect. This is not to say that credit markets operate perfectly in the region. If they did, then there would be no difference in outcomes between the information-only treatment and the treatment that included the subsidy. Rather, our results demonstrate that while liquidity is an issue, households are still able to take advantage of extension advice if it has been adapted to their context. This result is surprising, because the agricultural development literature frequently argues for the primacy of liquidity constraints and finds little evidence for the effectiveness of information-only interventions (Holden and Lunduka, 2013; Jones and Kondylis, 2018). Two important factors should be considered when comparing our results to those in the literature. First, Nigeria is among the countries in Africa with the highest levels of fertilizer use and fertilizer use is relatively common in Kano (Liverpool-Tasie, 2014). Second, previous studies of information-only interventions focus on the delivery of broad, standardized advice. The information provided by the DST in our study is tailored to each treated household. Thus, while the evidence regarding liquidity may lack external validity, the impact of digitally delivered personalized extension advice is likely to be generalizable to other settings.

4.2.Treatment effect on profit

From an economic point of view, the positive impact of the information-only treatment on yield cannot alone justify promotion of the app for scaling. Accordingly, in this subsection we focus on the effect of treatment on profit. Panel A of Table 3 presents the ITT effects of RiceAdvice on rice profit per hectare in the full set of data (both 2016 and 2017 outcomes). Panel B presents results from the initial year while Panel C presents results from the follow-up year.

We find clear evidence that personalized advice on nutrient management increases the profit from rice production. As with yields, the differences between OLS and ANCOVA estimates, as well as the differences with and without covariates, are minor. Households randomly assigned to the information-only treatment increase their profit by about $120, or 10 percent over control households. Gains were substantially larger for households that received the subsidized fertilizer, as they were likely able to reallocate funds to other productive activities. Households who received fertilizer in addition to RiceAdvice saw profit rise by about $275 or 23 percent compared to control households. As with the results for yields, there are significant differences between outcomes for those in each treatment. What is interesting, though, is that the fertilizer subsidy is not a necessary condition for households to make use of the recommendations from RiceAdvice.

To understand the effects of the intervention over time, we again estimate the ITT for each year separately. Profit for treated households is significantly higher than the control in each year. While the combined treatment effect is significant in each year, we again see a loss of significance for the information-only treatment effect in the second year. Yet, on average, the relative size of the gains made by households who received the information-only treatment do not change. Profits for those in the information-only group rose by 11 percent in 2016 and 9.5 percent in 2017. Because of this, we believe that the lack of significance reflects a lack of precision in the estimates and not a true null effect.

These findings imply that additional production costs related to the personalized advice are less than the gain in yield. While the size of the gains are subject to year-to-year variation in input and output prices, the gains are always positive. Again, our results contrast with much of the existing literature on the impact of information-only treatments. Indicative of this literature is Duflo et al. (2008), who find that blanket fertilizer advice from an official extension agency in Kenya has no effect on farm profits. We believe that such information-only interventions have been ineffective because the information provided is too general and fails to account for heterogeneity at the farm-level. By contrast, our intervention relies on a mobile DST that provides nutrient management advice adapted to the needs of the specific household.

5. Results for secondary outcomes

To investigate possible causal channels through which the adoption of personalized advice may influence yield and profit, we test the treatment effect on two intermediate outcomes: the quantity of fertilizer and the timing of fertilizer application.

5.1. Treatment effect on fertilizer quantity

The provision of personalized advice to farming households may increase or decrease the quantity of fertilizer depending on the initial distance of production to the efficiency frontier. Table 4 presents results from estimations of the ITT effect on the quantity of fertilizer used. Contrary to our priors, we find little evidence that personalized advice on nutrient management has an effect on the quantity of fertilizer. While the combined effect and the subsidy effect are significant, the preponderance of evidence is that those in the treatment groups use about the same amount of fertilizer as those in the control group.

There are two possible explanations for the mostly null results. First, it may be that some households increase their fertilizer quantity while other households decrease it, which would lead to a null effect on average. To investigate this explanation, we use a quantile regression, similar to Hossain et al. (2019), in order to estimate the slope and shape of the conditional distribution. We estimate the ITT effect on fertilizer quantity for three quartiles: the lower quartile (25^{th} percentile), the median quartile (50^{th} percentile), and the upper quartile (75^{th} percentile). We do not find substantial variation between quartiles in any of the treatment groups (see Table A.7 in the Appendix).

Second, rice farmers in the survey areas primarily use two types of fertilizer: NPK 15-15-15 and urea 46-0-0. The blanket extension advice provided by the Ministry of Agriculture and Rural Development breaks down fertilizer by deficiency in either nitrogen (N), phosphorous (P), and potassium (K). So, for both blanket advice and RiceAdvice, recommendations are for a specific compound fertilizer. Thus, the adoption of personalized extension advice may increase the quantity of one type of fertilizer while decreasing the quantity of the other type of fertilizer, which would result in a null effect on average. To investigate this second possible explanation, we model the treatment effect on NPK and urea quantities separately (see Table 5 and Table 6).

Our results show that personalized advice does have an effect on both the type and quantity of fertilizer in use – a result masked by a focus on average effects. Households in both treatment groups tend to reduce the amount of NPK and maintain or even increase the amount of urea, though the decrease in NPK (16-38 kg/ha) is larger than the increase in urea (7-12 kg/ha). We also find differences in NPK and urea use based on treatment arm. For NPK, the decrease in the amount of fertilizer use is similar for those in both the information-only and information plus subsidy treatments. When we directly compare T2 to T1, we find the differences in NPK use are not significant. For urea, the increase in quantity is driven by households in the information-only treatment. Households who received the recommended amount of urea fertilizer apply the same amount or slightly less urea as households in the control.

The NPK and urea regressions demonstrate that while the average amount of fertilizer used by the treatment groups was unaffected by the treatment, households did change their amount of fertilizer, just in offsetting ways. These results support Suri's (2011) conjecture that a focus on average effects masks highly heterogeneous returns to agricultural technologies. The blanket advice offered by extension agents may be correct for the average household, but no single household is exactly average. Relying on blanket recommendations, households may end up over-using one type of fertilizer while under-using a different type of fertilizer. Households provided with personalized extension advice adjust their application rates up or down, as needed. The result is a null effect on average, though households make adjustments to the quantity of each type of fertilizer they use.

5.2. Treatment effect on the application timing of fertilizer

In addition to the size of the fertilizer dose, application timing is vital to productive crop growth. For rice, the recommendation is to apply fertilizer four times during the growing season: basal (at transplanting or 16-20 days after sowing for direct seeding), tilling (36-40 days after sowing), panicle initiation (53-57 days after sowing), and booting (78-82 days after sowing). In Table 7 and Table 8, we assess the effect of personalized advice from RiceAdvice on the application timing of NPK and urea, respectively.

We find evidence of a negative effect on the application period (number of days after sowing) of NPK.[8] Treated households applied NPK about three days earlier than control households. As one would expect, there is no difference in application timing between the two treatment groups, since the only difference between the two treatments is the fertilizer subsidy.

[8] Only one application of each type of fertilizer is common among farmers. So the application time used is the number of days after sowing for the first application of fertilizer.

We find almost no evidence that the treatment had an effect on the application period of urea fertilizer, meaning that the treated and control households applied urea at approximately the same time after the sowing date. The differences between the treatment's effect on NPK and urea timing may be due to differences in familiarity with the type of fertilizer. Alternatively, it could be due to the different application period of each fertilizer, with NPK typically being applied at transplanting and urea typically being applied at tilling. However, even where differences are significant, the size of the effect is only two or three days, which may not be meaningful in an agronomic sense, since the application windows are five days in length.

5.3. Robustness and Limitations

While we find consistent and significant impacts of both the information-only treatment and the subsidy treatment on yields and profits, one may wonder if the results are robust to different inference, different samples, and different transformations of the data.

Our first check is to verify that the results are robust to adjustments for multiple comparisons. We calculate adjusted *p*-values as well as sharpened *q*-values following Clarke et al. (2019) and Anderson (2008), respectively. Results are reported in Table A.5 and Table A.6 in the Appendix. Across all regressions and using both methods, results that were significant remain significant when adjusting for multiple comparisons. The only exception is the timing of NPK application, where *p*-values and sharpened *q*-values remain small but adjust *p*-values become larger than 0.10.

Our second check is to verify that the results are robust to the missing data from households that chose not to cultivate rice in a given year. The number of missing observations is extremely small relative to the overall sample (see Table A.2 in the Appendix). We calculate Lee bounds on each treatment effect (Lee, 2009) and results are reported in Table A.8 in the Appendix. Our primary outcomes are bounded away from zero. In fact, for all outcomes where we saw a consistent treatment effect remain significant when accounting for the missing data.

Third, we implement a randomization inference procedure proposed by Heß (2017). In classical inference, the assumption is that the treatment is fixed and the sample is a random draw. In randomization inference, the assumption is that the sample is fixed and the assignment to treatment is a random draw. For each ANCOVA regression with covariates, we randomly permute the treatment indicator 5,000 times, taking into account the clustered design of the experiment. The procedure allows us to build a reference sample under the null hypothesis of no treatment effect. We then plot this distribution of outcomes when the hypothetical treatment effect is zero and compare it to the observed treatment effect. Figures A.2 and A.3 in the Appendix present these results. As with the other adjustments, results that are significant in the main analysis remain significant.

Next, we verify that our results are not due to a lack of balance in rice area across treatment groups. Our primary outcome variables are all in per hectare terms and so are adjusted for rice area post-treatment. Additionally, rice area is controlled for in the ANCOVA results because the per hectare outcomes pre-treatment are included as controls. Despite this, one may be concerned that impacts of yield or profit per hectare do not translate to total harvest and profit by affecting the area planted to rice. Table A.9 in the Appendix reports results from regressions in which total harvest, total profit, or rice area are the outcome. For harvest and profit, rice area in the baseline is include with the other covariates. In all cases for harvest and profit both treatments have a positive and significant impact on the two outcomes. Conversely, rice area is never significantly affected by either treatment.

Finally, we convert our primary outcomes from levels to logs. (see Table A.10 through Table A.16). We do this by using the inverse hyperbolic sine transformation, which allows us to retain zeros in the data as zeros, though the transformation only approximates the semi-elasticities from standard log transforms (Burbidge et al., 1988; Bellemare and Wichman, 2020). Of the 84 regressions, results change from not significant as levels to significant as logs in 16 regressions. In four regressions results that were significant as levels are no longer significant as logs. The transformation does not change the significance (or lack of significance) for coefficient estimates in the remaining 60 regressions. The greatest sensitivity to the transformation are results regarding quantity of NPK or urea fertilizer. While both treatments significantly reduced the quantity of NPK applied over both years, this result is not robust when quantities are logged. Conversely, coefficients on T2 for 2017 for NPK application move from not significant to significant as do coefficients on T1 for 2016 for urea application. Of the 20 changes in significance, 12 occur in NPK or urea application rates. We conclude that there is some sensitivity regarding estimates when fertilizer is broken down by type.

While our main results are robust to a number of alternative methods, the limitations of our study should be noted. First, because our study is not a 2x2 design, we are unable to measure the effect of only the subsidy or the overall effect of RiceAdvice. We are only able to directly estimate the impact of the information-only treatment and the impact of RiceAdvice combined with the fertilizer subsidy. We can calculate the total effect of both treatments and the differences between treatments but this requires a linearity assumption that may not hold. Second, while we believe that RiceAdvice is effective because it provides personalized recommendations that account for heterogeneity in the soil, we are unable to directly test this. We lack detailed data on soil and thus cannot directly measure if soil heterogeneity is the primary mechanism through which the app works. Finally, like all RCTs, the results can only be said to hold for those who are represented by study participants. In our case, rice growing households in Kano, Nigeria. While we have no reason to suspect that RiceAdvice will be ineffective with rice growing households in other parts of the world, our study is unable to demonstrate this conjecture.

6. Conclusion

Historically, governments and development agencies have found it necessary to use abstracted, simplified, blank advice when promoting the adoption of improved technologies. One potential reason why adoption of these technologies has remained low is that while on average the recommendations are accurate, for any given household the recommendations will be wrong. The local environment at any particular farm means that blanket advice on fertilizer application will result in the farmer over or under using the input, with negative consequences for yield and profit. Until recently, the cost of adapting extension advice to local conditions was prohibitively expensive. But, with advances in mobile technology, decision support tools (DSTs) can be developed and disseminated at greatly reduced cost. By using DSTs, farmers and extension agents can fine-tune management practices by taking into account variations in local environmental and economic conditions, reduce their inefficiencies, and shorten the learning process.

We explore the potential for an Android-based DST called RiceAdvice. The mobile app allows households to provide information to extension agents about their local growing conditions, production costs, and market information. The extension agent can then use the app to provide recommendations for a nutrient management plan designed to increase both yields and profits. Using a randomized control trial, we find that households with access to RiceAdvice increase their yields and profits. We also find that households are able to take advantage of the personalized extension information within their current credit constraints. Households in the information-only treatment arm are still able to significantly improve their outcomes, though not by as much as households who received a subsidy to cover the full cost of the recommended fertilizer amount.

These outcomes are not driven by an overall increase in the use of fertilizer. On average there is no significant difference between fertilizer application rates for treatment and control households. Rather, the personalized extension advice allows households that previously over-used fertilizer to reduce their application rate and households who previously under-used fertilizer to increase their application rate. The study resulted in households increasing their yields and income while having a net zero effect on the total amount of fertilizer used. This suggests that improvements to productivity and livelihoods need not come at the cost of increased overall chemical fertilizer use, and the corresponding negative effects on the environment.

While our results are specific to a particular DST, they add to a nascent body of literature suggesting that the null results typical of information-only interventions may be due to the overly broad information provided in the studies. In the case of technology adoption, how individuals are taught to use the technology plays an important role in whether that technology is beneficial or not. For farming households looking to take advantage of new seeds and other improved inputs, the revolution in mobile technology allows for a move away from the old one-size-fits-all advice and towards the delivery of personalized and profitable recommendations.

References

Abay, K.A., Abate, G.T., Barrett, C.B., and Bernard, T. 2019. Correlated non-classical measurement errors, `Second best' policy inferences, and the inverse size-productivity relationship in agriculture. *Journal of Development Economics 139*: 171-184.

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 103* (484): 1481–1495.

Ashraf, N., Jack, B.K., and Kamenica, E. 2013. Information and subsidies: Complements or substitutes? *Journal of Economic Behavior & Organization 88*: 133–139.

Arouna, A., Michler, J. and Lokossou, J. 2019. Contract farming and rural transformation: Evidence from a field experiment in Benin. NBER Working Paper No. w25665.

Bellemare, M.F. 2010. Agricultural extension and imperfect supervision in contract farming: Evidence from Madagascar. *Agricultural Economics 41*: 507-517.

Bellemare, M. F., & Wichman, C. J. (2020). Elasticities and the inverse hyperbolic sine transformation. *Oxford Bulletin of Economics and Statistics*, 82(1), 50-61.

Bettinger, E.P., Long, B.T., Oreopoulos, P. and Sanbonmatsu, L. 2012. The role of application assistance and information in college decisions: Results from the H&R Block FAFSA experiment. *Quarterly Journal of Economics 127* (3):1205–1242.

Bryan, G., Chowdhury, S., and Mobarak, A.M. 2014. Underinvestment in a profitable technology: The case of seasonal migration in Bangladesh. *Econometrica 82* (5):1671–1748.

Burbidge, J.B., Magee, L., and Robb, A.L. 1988. Alternative transformations to handle extreme values of the dependent variable. *Journal of the American Statistical Association 83*: 123-7.

Burchardi, K.B., Gulesci, S., Lerva, B., and Sulaim, M. 2019. Moral hazard: Experimental evidence from tenancy contracts. *Quarterly Journal of Economics 134* (1): 281–347.

Byerlee, D., and Deininger, K. 2013. Growing resource scarcity and global farmland investment. *Annual Review of Resource Economics 5* (1): 13–34.

Cameron, A.C., and Miller, D.L. 2015. A practitioner's guide to cluster-robust inference. *Journal of Human Resources 50*: 317-372.

Casaburi, L., and Willis, J. 2018. Time versus state in insurance: Experimental evidence from contract farming in Kenya. *American Economic Review 108* (12): 3778–3813.

Casaburi, L., Kremer, M., Mullainathan, S., and Ramrattan, R. 2019. Harnessing ICT to increase agricultural production: Evidence from Kenya. University of Zurich, mimeo.

Chude, V.O., Olayiwola, S.O., Daudu, C., Ekeoma, A., eds. 2012. *Fertilizer Use and Management Practices for Crops in Nigeria*. 4th ed. Abuja: Federal Ministry of Agriculture and Rural Development.

Clarke, D., Romano, J.P., and Wolf, M. 2019. The Romano-Wolf multiple hypothesis correction in Stata. IZA Discussion Paper Series No. 12845.

Cole, S.A., and Fernando, N. 2018. 'Mobile'izing agricultural advice: Technology adoption, diffusion and sustainability. Harvard University, mimeo.

Conley, T.G., and Udry, C.R. 2010. Learning about a new technology: Pineapple in Ghana. *American Economic Review 100* (1): 35–69.

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 98* (2): 482–488.

Fafchamps, M., and Minten, B. 2012. Impact of SMS-based agricultural information on Indian farmers. *World Bank Economic Review 26* (3): 383–414.

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

Gourlay, S., Kilic, T., and Lobell, D.B. 2019. A new spin on an old debate: Errors in farmer-reported production and their implications for inverse scale - Productivity relationship in Uganda. *Journal of Development Economics 141*: 102376.

Grêt-Regamy, A., Sirén, E., Brunner, S.H., and Weibel, B. 2017. Review of decision support tools to operationalize the ecosystem services concept. *Ecosystem Services 26* (B): 306–315.

Heckman, J.J., Ichimura, H., and Todd, P.E. 1998. Matching as an econometric evaluation estimator. *Review of Economic Studies 65* (2): 261–294.

Heß, S.H. 2017. Randomization inference with Stata: A guide and software. *Stata Journal 17* (3): 630-651.

Holden, S.T., and Lunduka, R.W. 2013. Input subsidies, cash constraints, and timing of input supply. *American Journal of Agricultural Economics 96* (1): 290–307.

Hossain, M., Malek, M.A., Hossain, M.A., Reza, M.H., and Ahmed, M.S. 2019. Agricultural microcredit for tenant farmers: Evidence from a field experiment in Bangladesh. *American Journal of Agricultural Economics 101* (3): 692-709.

Imai, K., King, G., and Nall, C. 2009. The essential role of pair matching in cluster-randomized experiments, with application to the Mexican universal health insurance evaluation. *Statistical Science 24* (1): 29–53.

Jack, B.K. 2011. Constraints on the adoption of agricultural technologies in developing countries. *White paper, Agricultural Technology Adoption Initiative*. Boston: J-PAL (MIT) and Berkeley: CEGA (UC Berkeley).

Jones, M. and Kondylis, F. 2018. Does feedback matter? Evidence from agricultural services. *Journal of Development Economics 131*: 28-41.

Karlan, D., Osei, R., Osei-Akoto, I., and Udry, C. 2014. Agricultural decisions after relaxing credit and risk constraints. *Quarterly Journal of Economics 129* (2): 597–652.

Kondylis, F., Mueller, V., Sheriff, G., and Zhu, S. 2016. Do female instructors reduce gender bias in diffusion of sustainable land management techniques? Experimental evidence from Mozambique. *World Development 78*: 436-449.

Kondylis, F., Mueller, V., and Zhu, J. 2017. Seeing is believing? Evidence from an extension network experiment. *Journal of Development Economics 125*: 1-20.

Krishnan, P., and Patnam, M. 2014. Neighbors and extension agents in Ethiopia: Who matters more for technology adoption? *American Journal of Agricultural Economics 96* (1): 308-327.

Ksoll, C., Lilleør, H.B., Lønborg, J.H., and Rasmussen, O.D. 2016. Impact of village savings and loan associations: Evidence from a cluster randomized trial. *Journal of Development Economics 120*: 70–85.

Lee, D. 2009. Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *Review of Economic Studies 76* (3): 1071-102.

Liu, E.M. 2013. Time to change what to sow: Risk preferences and technology adoption decisions of cotton farmers in China. *Review of Economics and Statistics 95* (4): 1386-1403.

Liverpool-Tasie, L.S.O. 2014. Fertilizer subsidies and private market participation: The case of Kano State, Nigeria. *Agricultural Economics 45*: 663–678.

MacCarthy, D.S., Kihara, J., Masikati, P., and Adiku, S.G.K. 2017. Decision support tools for site-specific fertilizer advice and agricultural planning in selected countries in sub-Sahara Africa. *Nutrient Cycling in Agroecosystems 110* (3): 343–359.

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

McKenzie, D. 2012. Beyond baseline and follow-up: The case for more T in experiments. *Journal of Development Economics 99*: 210-21.

Michler, J.D., Tjernström, E., Verkaart, S., and Mausch, K. 2019. Money matters: The role of yields and profits in agricultural technology adoption. *American Journal of Agricultural Economics 101* (3): 710–731.

Mullally, C. and Maffioli, A. 2016. Extension and matching grants for improved management: An evaluation of the Uruguayan Livestock Program. *American Journal of Agricultural Economics 98* (1): 333-350.

Pan, Y., Smith, S.C., and Sulaiman, M. 2018. Agricultural extension and technology adoption for food security: Evidence from Uganda. *American Journal of Agricultural Economics 100* (4): 1012-1031.

Rose, D.C., Sutherland, W.J., Parker, C., Lobley, M., Winter, M., Morris, C., Twining, S., Ffoulkes, C., Amano, T., and Dicks, L.V. 2016. Decision support tools for agriculture: Towards effective design and delivery. *Agricultural Systems 149*: 165–174.

Saito, K. Diack, S., Dieng, I., and Ndiaye, M.K. 2015a. On-farm testing of a nutrient management decision-support tool for rice in the Senegal River valley. *Computers and Electronics in Agriculture 116*: 36–44.

Saito, K., Dieng, I., Toure, A., Somado, E.A., and Wopereis, M.C.S. 2015b. Rice yield growth analysis for 24 African countries over 1960-2012. *Global Food Security 5*: 62–69.

Scott, J.C. 1998. *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. New Haven: Yale University Press.

Suri, T. 2011. Selection and comparative advantage in technology adoption. *Econometrica 79* (1): 159–209.

Tjernström, E., Lybbert, T.J., Frattarola Hernández, R., and Correa, J.S. 2020. Learning by (virtually) doing: Experimentation and belief updating in smallholder agriculture. University of Sydney, mimeo.

USDA, 2018. PSD Online custom query. http://apps.fas.usda.gov/psdonline/psdQuery.aspx. Accessed 12 December 2018.

Van Campenhout, B., Spielman, D.J., and Lecoutere, E. Forthcoming. Information and communication technologies to provide agricultural advice to smallholder farmers: Experimental evidence from Uganda. *American Journal of Agricultural Economics*.

van Oort, P.A.J., Saito, K., Dieng, I., Grassini, P., Cassman., K.G., and van Ittersum, M.K., 2017. Can yield gap analysis be used to inform R&D prioritisation? *Global Food Security 12*: 109–118.

Walz, S.P., and Deterding, S. 2015. *The Gameful World: Approaches, Issues, Applications*. Cambridge: MIT Press.

Watkiss, P., Hunt, A., Blyth, W., and Dyszynski, J. 2015. The use of new economic decision support tools for adaptation assessment: A review of methods and applications, towards guidance on applicability. *Climate Change 132*: 401–416.

Figure 1.A: Screenshot of RiceAdvice data inputs

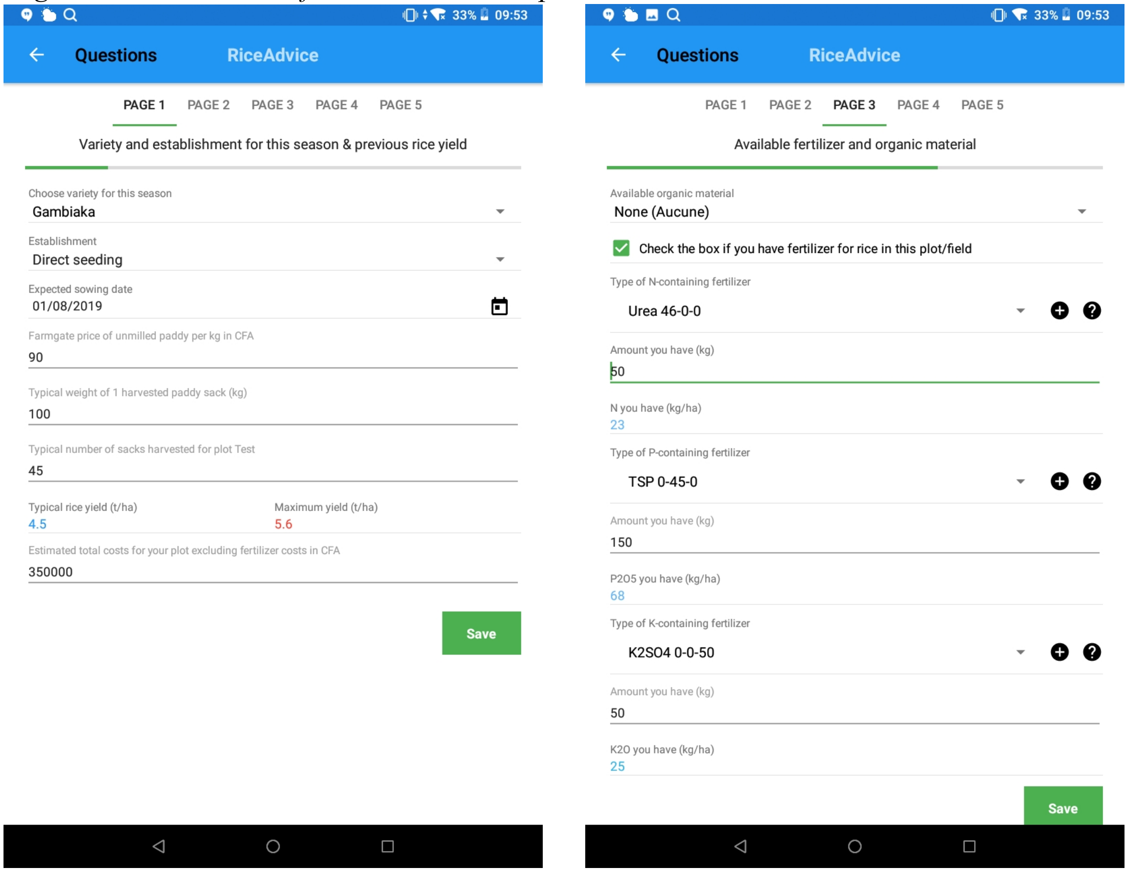

Figure 1.B: Screenshot of RiceAdvice outputs

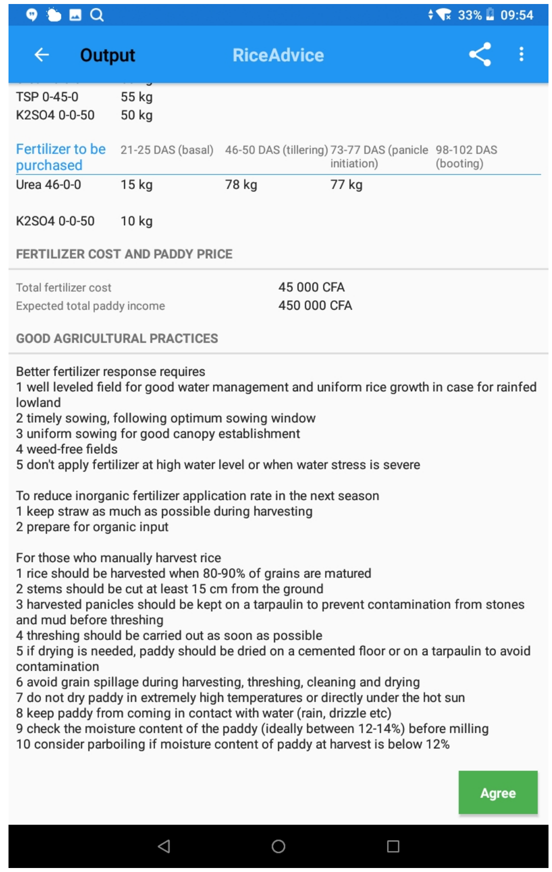

Figure 2: Experiment design

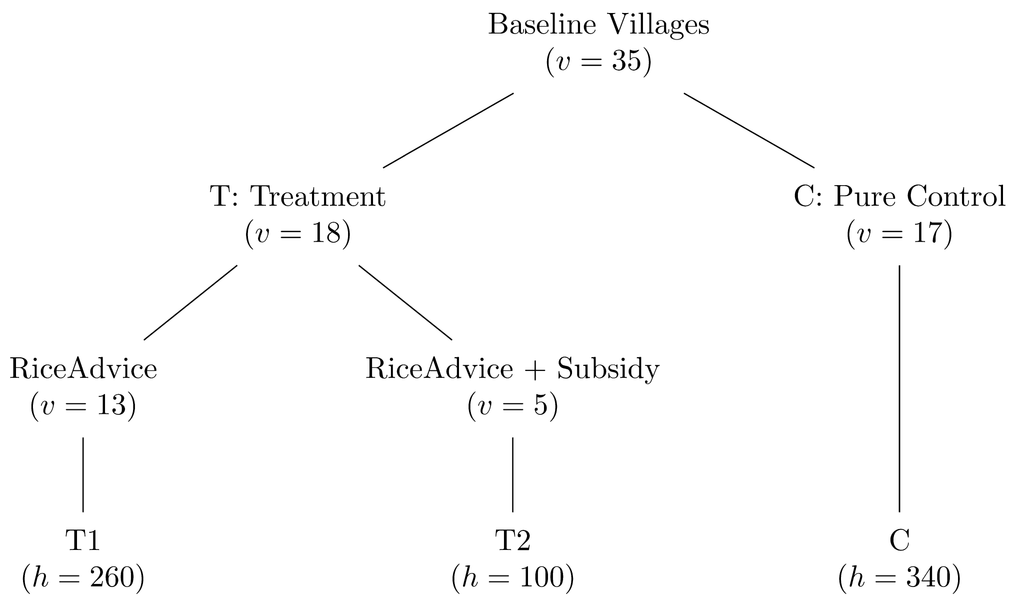

Figure 3: Outcomes by treatment group

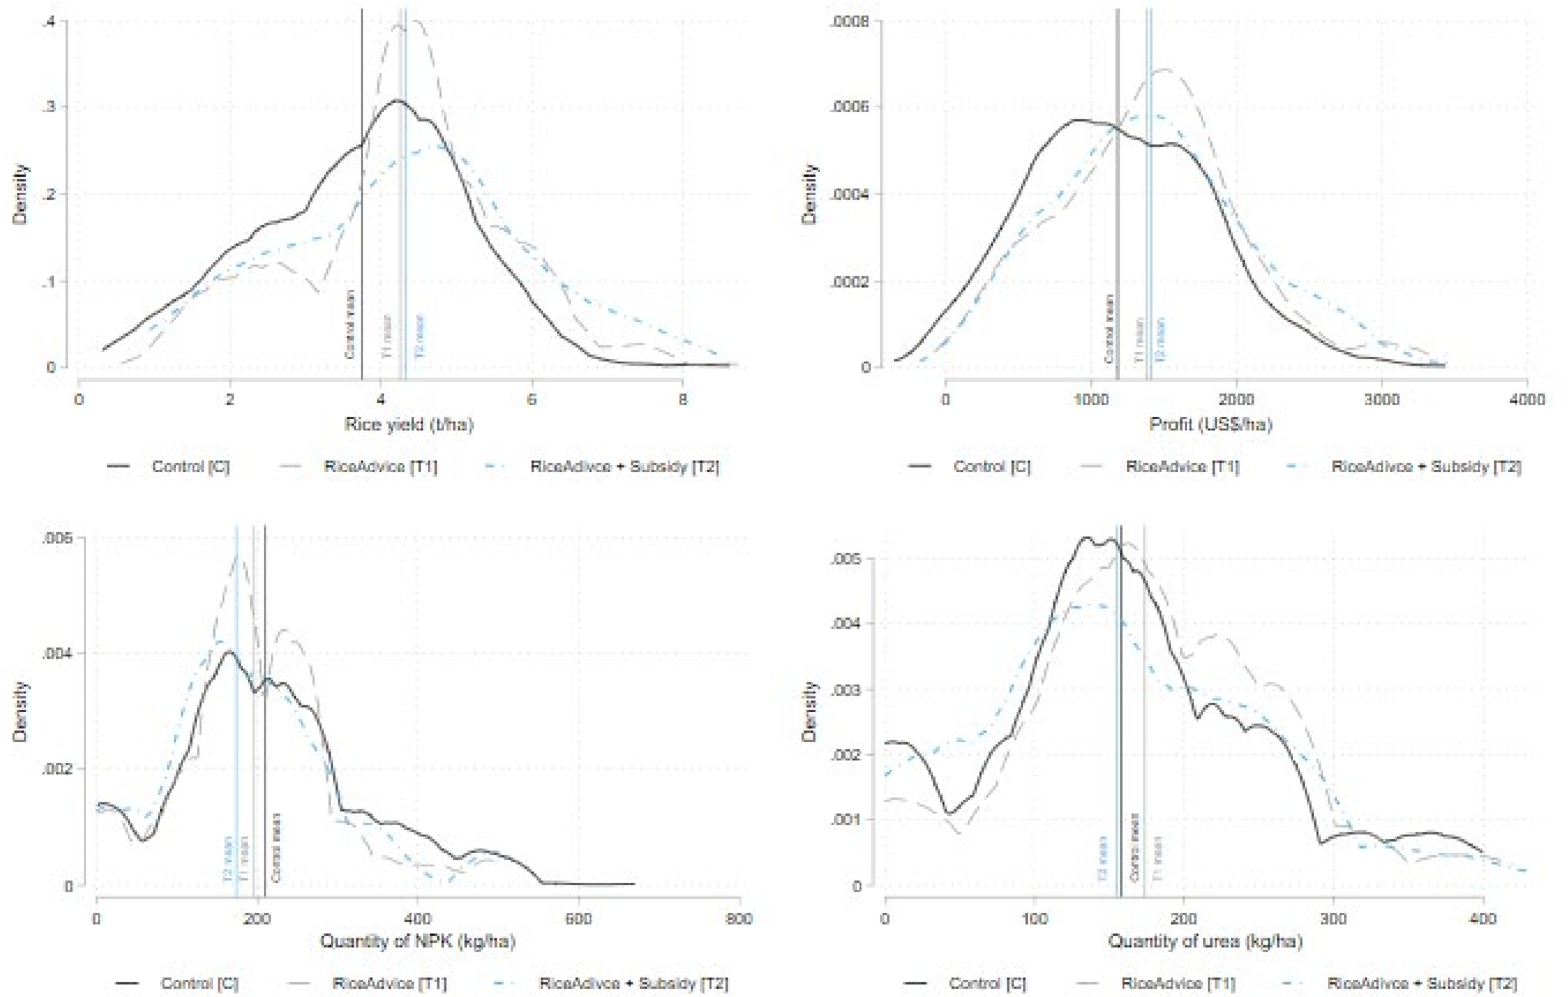

Note: The figures show the distribution of post-experiment (2016 and 2017 pooled) values for each outcome variable by treatment group. Vertical lines mark the mean value for each outcome variable by treatment group.

Table 1: Baseline characteristics and balance pre-contamination

	Control group [C] (1)	Difference with Treated [T-C] (2)	Difference with T1 [T1-C] (3)	Difference with T2 [T2-C] (4)
Household characteristics				
Age of rice farmer (year)	37.47 (11.27)	-1.631	-2.372*	0.237
Household size (n)	11.52 (7.670)	0.319	0.507	-0.082
Formal education (=1)	0.256 (0.437)	-0.005	0.014	-0.045
Farming is main activity (=1)	0.879 (0.326)	-0.008	-0.022	0.032
Agricultural training days (n)	0.653 (2.404)	0.270	0.082	0.834
Access to credit (=1)	0.144 (0.352)	0.011	0.007	0.025
Rice area (ha)	0.740 (0.508)	0.237***	0.225**	0.296*
Production values				
Rice yield (t/ha)	3.484 (1.759)	-0.189	-0.242	-0.062
Rice income (US$/ha)	1,675 (845.9)	-91.02	-116.6	-29.62
Profit (US$/ha)	1,357 (797.6)	-83.92	-106.8	-30.54
Quantity of NPK (kg/ha)	184.0 (87.71)	-3.250	-6.989	6.380
Quantity of urea (kg/ha)	164.0 (88.91)	10.57	4.734	25.64***

Note: The first column reports means of the data in the control group at baseline with standard deviations in parentheses. Columns (2) - (4) report coefficients from OLS regressions of the variables of interest on treatment status within different groups and represent the difference of treatment minus control. Significance tests are based on standard errors clustered at the village-level (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table 2: Treatment effects on rice yield (t/ha)

	OLS (1)	OLS (2)	ANCOVA (3)	ANCOVA (4)
Panel A: all years				
RiceAdvice [T1]	0.253**	0.249**	0.260**	0.258**
	(0.123)	(0.116)	(0.125)	(0.118)
RiceAdvice + Subsidy [T2]	0.737***	0.725***	0.736***	0.728***
	(0.125)	(0.117)	(0.127)	(0.120)
Combined treatment [T]	0.990***	0.974***	0.996***	0.986***
	(0.222)	(0.211)	(0.229)	(0.219)
Difference between treatments [T2-T1]	0.484***	0.477***	0.476***	0.470***
	(0.109)	(0.098)	(0.105)	(0.095)
Mean dependent variable in control	3.755			
Observations	1,368	1,368	1,353	1,353
R-squared	0.214	0.221	0.215	0.222
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes
Panel B: 2016 only				
RiceAdvice [T1]	0.252**	0.252**	0.282**	0.282**
	(0.115)	(0.111)	(0.124)	(0.120)
RiceAdvice + Subsidy [T2]	0.619***	0.611***	0.632***	0.628***
	(0.123)	(0.123)	(0.131)	(0.132)
Combined treatment [T]	0.872***	0.862***	0.914***	0.910***
	(0.222)	(0.218)	(0.240)	(0.237)
Difference between treatments [T2-T1]	0.367***	0.359***	0.350***	0.346***
	(0.088)	(0.088)	(0.086)	(0.088)
Mean dependent variable in control	3.782			
Observations	694	694	686	686
R-squared	0.227	0.230	0.234	0.237
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes
Panel C: 2017 only				
RiceAdvice [T1]	0.247	0.237	0.234	0.227
	(0.179)	(0.172)	(0.180)	(0.174)
RiceAdvice + Subsidy [T2]	0.847***	0.827***	0.833***	0.817***
	(0.198)	(0.177)	(0.202)	(0.182)
Combined treatment [T]	1.094***	1.064***	1.067***	1.044***
	(0.323)	(0.306)	(0.330)	(0.314)
Difference between treatments [T2-T1]	0.599***	0.591***	0.599***	0.590***
	(0.196)	(0.168)	(0.195)	(0.168)
Mean dependent variable in control	3.728			
Observations	674	674	667	667
R-squared	0.320	0.336	0.320	0.336
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes

Note: For simplicity, coefficient estimates are only reported for the treatment effect. All regressions include LGA (strata) fixed effects. ANCOVA estimates include baseline yields and year dummies, where applicable. Covariates include household size, age of household head, number of days in agricultural training, and indicators for if the household head has formal education, if farming is the household's main activity, and if they have access to credit. Standard errors clustered at the village-level are in parentheses (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table 3: Treatment effects on rice profit (US$/ha)

	OLS (1)	OLS (2)	ANCOVA (3)	ANCOVA (4)
Panel A: all years				
RiceAdvice [T1]	115.6**	118.9**	122.0**	126.1**
	(50.69)	(48.95)	(51.29)	(49.58)
RiceAdvice + Subsidy [T2]	275.9***	273.2***	282.0***	279.9***
	(50.77)	(48.22)	(51.00)	(48.94)
Combined treatment [T]	391.5***	392.1***	404.0***	406.1***
	(92.85)	(89.24)	(94.19)	(90.90)
Difference between treatments [T2-T1]	160.2***	154.3***	160.0***	153.8***
	(40.90)	(38.44)	(39.88)	(38.01)
Mean dependent variable in control	1,181			
Observations	1,368	1,368	1,353	1,353
R-squared	0.14	0.15	0.33	0.33
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes
Panel B: 2016 only				
RiceAdvice [T1]	153.9**	153.9**	164.7**	166.0**
	(61.35)	(62.28)	(66.05)	(67.10)
RiceAdvice + Subsidy [T2]	259.0***	261.7***	264.6***	268.4***
	(67.14)	(66.34)	(69.39)	(68.98)
Combined treatment [T]	412.9***	415.6***	429.3***	434.4***
	(118.0)	(116.5)	(125.7)	(124.6)
Difference between treatments [T2-T1]	105.1**	107.8*	99.82*	102.4*
	(51.30)	(54.57)	(50.46)	(054.71)
Mean dependent variable in control	1,427			
Observations	694	694	686	686
R-squared	0.18	0.18	0.18	0.19
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes
Panel C: 2017 only				
RiceAdvice [T1]	82.54	88.30	76.68	82.90
	(56.76)	(53.74)	(57.52)	(54.79)
RiceAdvice + Subsidy [T2]	301.4***	292.2***	296.0***	286.9***
	(68.10)	(60.63)	(69.65)	(62.40)
Combined treatment [T]	383.9***	380.5***	372.7***	369.8***
	(105.9)	(98.70)	(108.4)	(101.6)
Difference between treatments [T2-T1]	218.84***	203.9***	219.3***	204.0***
	(67.05)	(58.19)	(67.52)	(58.90)
Mean dependent variable in control	923.4			
Observations	674	674	667	667
R-squared	0.28	0.30	0.28	0.29
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes

Note: For simplicity, coefficient estimates are only reported for the treatment effect. All regressions include LGA (strata) fixed effects. ANCOVA estimates include baseline profits and year dummies, where applicable. Covariates include household size, age of household head, number of days in agricultural training, and indicators for if the household head has formal education, if farming is the household's main activity, and if they have access to credit. Standard errors clustered at the village-level are in parentheses (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table 4: Treatment effects on quantity of fertilizer (kg/ha)

	OLS	OLS	ANCOVA	ANCOVA
	(1)	(2)	(3)	(4)
Panel A: all years				
RiceAdvice [T1]	-8.777	-11.775	-8.179	-10.92
	(10.61)	(10.08)	(10.55)	(10.01)
RiceAdvice + Subsidy [T2]	-31.67***	-33.62***	-31.00***	-32.65***
	(10.98)	(11.45)	(10.98)	(11.40)
Combined treatment [T]	-40.45**	-45.39**	-39.18**	-43.57**
	(18.69)	(19.25)	(18.67)	(19.19)
Difference between treatments [T2-T1]	-22.90**	-21.84**	-22.82**	-21.74**
	(10.81)	(9.738)	(10.73)	(9.601)
Mean dependent variable in control	366.6			
Observations	1,368	1,368	1,353	1,353
R-squared	0.068	0.080	0.100	0.113
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes
Panel B: 2016 only				
RiceAdvice [T1]	4.156	2.064	3.894	1.694
	(14.82)	(15.45)	(15.15)	(15.75)
RiceAdvice + Subsidy [T2]	-4.215	-10.18	-5.719	-11.63
	(30.22)	(29.85)	(30.25)	(29.91)
Combined treatment [T]	-0.059	-8.114	-1.825	-9.936
	(38.93)	(39.33)	(39.31)	(39.72)
Difference between treatments [T2-T1]	-8.371	-12.24	-9.614	-13.32
	(27.40)	(26.68)	(27.25)	(26.61)
Mean dependent variable in control	316.9			
Observations	694	694	686	686
R-squared	0.076	0.099	0.078	0.102
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes
Panel C: 2017 only				
RiceAdvice [T1]	-23.05	-26.94	-20.67	-23.94
	(17.90)	(16.26)	(17.60)	(15.73)
RiceAdvice + Subsidy [T2]	-61.04***	-59.08***	-56.89***	-54.35***
	(18.65)	(16.44)	(18.36)	(15.94)
Combined treatment [T]	-84.08**	-86.02***	-77.56**	-78.29***
	(31.51)	(28.52)	(30.49)	(27.11)
Difference between treatments [T2-T1]	-37.99**	-32.14*	-36.22*	-30.41*
	(18.53)	(16.00)	(19.08)	(16.36)
Mean dependent variable in control	418.4			
Observations	674	674	667	667
R-squared	0.088	0.115	0.091	0.119
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes

Note: For simplicity, coefficient estimates are only reported for the treatment effect. All regressions include LGA (strata) fixed effects. ANCOVA estimates include baseline fertilize use and year dummies, where applicable. Covariates include household size, age of household head, number of days in agricultural training, and indicators for if the household head has formal education, if farming is the household's main activity, and if they have access to credit. Standard errors clustered at the village-level are in parentheses (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table 5: Treatment effects on quantity of NPK fertilizer (kg/ha)

	OLS	OLS	ANCOVA	ANCOVA
	(1)	(2)	(3)	(4)
Panel A: all years				
RiceAdvice [T1]	-14.07**	-16.06**	-14.22**	-16.03**
	(6.142)	(6.195)	(6.204)	(6.272)
RiceAdvice + Subsidy [T2]	-21.12***	-21.81***	-21.48***	-21.98***
	(7.038)	(7.106)	(7.059)	(7.151)
Combined treatment [T]	-35.20***	-37.87***	-35.70***	-38.01***
	(11.88)	(12.13)	(11.98)	(12.25)
Difference between treatments [T2-T1]	-7.047	-5.750	-7.266	-5.945
	(5.784)	(5.543)	(5.758)	(5.548)
Mean dependent variable in control	210.4			
Observations	1,368	1,368	1,353	1,353
R-squared	0.049	0.057	0.131	0.140
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes
Panel B: 2016 only				
RiceAdvice [T1]	-12.44	-13.25	-12.07	-13.03
	(9.161)	(9.764)	(9.304)	(9.891)
RiceAdvice + Subsidy [T2]	-21.99	-24.90	-21.80	-24.68
	(19.00)	(18.16)	(18.95)	(18.15)
Combined treatment [T]	-34.43	-38.15	-33.88	-37.71
	(24.24)	(23.93)	(24.38)	(24.09)
Difference between treatments [T2-T1]	-9.543	-11.66	-9.734	-11.65
	(17.38)	(16.66)	(17.22)	(16.56)
Mean dependent variable in control	176.8			
Observations	694	694	686	686
R-squared	0.056	0.075	0.057	0.076
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes
Panel C: 2017 only				
RiceAdvice [T1]	-16.97*	-20.10**	-16.68*	-19.32**
	(9.713)	(9.258)	(9.704)	(9.122)
RiceAdvice + Subsidy [T2]	-22.19	-20.72	-21.78	-19.91
	(14.67)	(12.48)	(14.59)	(12.50)
Combined treatment [T]	-39.16*	-40.82**	-38.45*	-39.24**
	(20.91)	(18.70)	(20.81)	(18.59)
Difference between treatments [T2-T1]	-5.220	-0.625	-5.104	-0.591
	(13.48)	(11.55)	(13.45)	(11.56)
Mean dependent variable in control	245.5			
Observations	674	674	667	667
R-squared	0.085	0.111	0.085	0.111
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes

Note: For simplicity, coefficient estimates are only reported for the treatment effect. All regressions include LGA (strata) fixed effects. ANCOVA estimates include baseline NPK use and year dummies, where applicable. Covariates include household size, age of household head, number of days in agricultural training, and indicators for if the household head has formal education, if farming is the household's main activity, and if they have access to credit. Standard errors clustered at the village-level are in parentheses (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table 6: Treatment effects on quantity of urea fertilizer (kg/ha)

	OLS	OLS	ANCOVA	ANCOVA
	(1)	(2)	(3)	(4)
Panel A: all years				
RiceAdvice [T1]	12.76**	11.95**	12.79**	12.10**
	(5.039)	(4.752)	(5.182)	(4.912)
RiceAdvice + Subsidy [T2]	-4.579	-5.818	-3.924	-4.912
	(5.634)	(5.589)	(5.886)	(5.857)
Combined treatment [T]	8.183	6.130	8.868	7.183
	(9.470)	(9.531)	(9.893)	(10.00)
Difference between treatments [T2-T1]	-17.34***	-17.77***	-16.72***	-17.01***
	(4.959)	(4.098)	(5.012)	(4.100)
Mean dependent variable in control	158.0			
Observations	1,368	1,368	1,353	1,353
R-squared	0.081	0.093	0.089	0.101
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes
Panel B: 2016 only				
RiceAdvice [T1]	16.60*	15.31	16.03	14.80
	(9.494)	(9.377)	(9.527)	(9.422)
RiceAdvice + Subsidy [T2]	17.77	14.73	16.49	13.55
	(14.53)	(14.94)	(14.62)	(15.05)
Combined treatment [T]	34.37	30.04	32.51	28.35
	(21.13)	(21.52)	(21.25)	(21.67)
Difference between treatments [T2-T1]	1.172	-0.586	0.461	-1.255
	(12.50)	(12.62)	(12.55)	(12.69)
Mean dependent variable in control	140.1			
Observations	694	694	686	686
R-squared	0.073	0.093	0.073	0.092
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes
Panel C: 2017 only				
RiceAdvice [T1]	8.407	7.960	9.157	8.884
	(7.250)	(6.639)	(7.392)	(6.874)
RiceAdvice + Subsidy [T2]	-27.76***	-27.26***	-24.88***	-23.97***
	(7.821)	(8.286)	(7.719)	(8.210)
Combined treatment [T]	-19.35	-19.30	-15.72	-15.08
	(13.12)	(12.85)	(13.06)	(12.93)
Difference between treatments [T2-T1]	-36.17***	-35.22***	-34.03***	-32.85***
	(7.446)	(7.776)	(7.608)	(7.877)
Mean dependent variable in control	176.7			
Observations	674	674	667	667
R-squared	0.115	0.136	0.127	0.150
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes

Note: For simplicity, coefficient estimates are only reported for the treatment effect. All regressions include LGA (strata) fixed effects. ANCOVA estimates include baseline urea use and year dummies, where applicable. Covariates include household size, age of household head, number of days in agricultural training, and indicators for if the household head has formal education, if farming is the household's main activity, and if they have access to credit. Standard errors clustered at the village-level are in parentheses (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table 7: Treatment effects on timing of first NPK application (days)

	OLS	OLS	ANCOVA	ANCOVA
	(1)	(2)	(3)	(4)
Panel A: all years				
RiceAdvice [T1]	-1.545***	-1.482***	-1.608***	-1.565***
	(0.465)	(0.485)	(0.499)	(0.521)
RiceAdvice + Subsidy [T2]	-1.308**	-1.226*	-1.293**	-1.235**
	(0.622)	(0.639)	(0.573)	(0.602)
Combined treatment [T]	-2.852***	-2.708***	-2.901***	-2.800***
	(0.915)	(0.950)	(0.907)	(0.956)
Difference between treatments [T2-T1]	0.237	0.256	0.315	0.330
	(0.608)	(0.620)	(0.577)	(0.596)
Mean dependent variable in control	15.40			
Observations	1,215	1,215	1,150	1,150
R-squared	0.165	0.174	0.214	0.224
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes
Panel B: 2016 only				
RiceAdvice [T1]	-1.121**	-1.049**	-1.168**	-1.120*
	(0.479)	(0.504)	(0.520)	(0.562)
RiceAdvice + Subsidy [T2]	-1.585***	-1.464**	-1.539***	-1.434***
	(0.496)	(0.537)	(0.456)	(0.510)
Combined treatment [T]	-2.706***	-2.513***	-2.707***	-2.554***
	(0.846)	(0.890)	(0.860)	(0.931)
Difference between treatments [T2-T1]	-0.464	-0.414	-0.371	-0.314
	(0.486)	(0.542)	(0.465)	(0.534)
Mean dependent variable in control	16.19			
Observations	606	606	576	576
R-squared	0.085	0.105	0.109	0.130
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes
Panel C: 2017 only				
RiceAdvice [T1]	-2.186***	-2.125***	-2.237***	-2.169***
	(0.653)	(0.674)	(0.709)	(0.728)
RiceAdvice + Subsidy [T2]	-1.187	-1.122	-1.177	-1.120
	(0.988)	(1.019)	(0.919)	(0.937)
Combined treatment [T]	-3.373**	-3.247**	-3.415**	-3.289**
	(1.417)	(1.466)	(1.414)	(1.453)
Difference between treatments [T2-T1]	1.000	1.003	1.060	-1.049
	(0.894)	(0.913)	(0.833)	(0.838)
Mean dependent variable in control	14.65			
Observations	609	609	574	574
R-squared	0.394	0.402	0.396	0.407
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes

Note: For simplicity, coefficient estimates are only reported for the treatment effect. All regressions include LGA (strata) fixed effects. ANCOVA estimates include baseline timing of NPK application and year dummies, where applicable. Covariates include household size, age of household head, number of days in agricultural training, and indicators for if the household head has formal education, if farming is the household's main activity, and if they have access to credit. Standard errors clustered at the village-level are in parentheses (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table 8: Treatment effects on timing of first urea application (days)

	OLS	OLS	ANCOVA	ANCOVA
	(1)	(2)	(3)	(4)
Panel A: all years				
RiceAdvice [T1]	-0.803	-0.805	-0.639	-0.595
	(0.653)	(0.679)	(0.652)	(0.691)
RiceAdvice + Subsidy [T2]	-0.397	-0.258	-0.589	-0.434
	(0.679)	(0.708)	(0.696)	(0.712)
Combined treatment [T]	-1.200	-1.063	-1.228	-1.029
	(1.275)	(1.330)	(1.257)	(1.308)
Difference between treatments [T2-T1]	0.406	0.547	0.050	0.162
	(0.386)	(0.398)	(0.487)	(0.506)
Mean dependent variable in control	33.03			
Observations	1,226	1,226	1,168	1,168
R-squared	0.103	0.112	0.114	0.123
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes
Panel B: 2016 only				
RiceAdvice [T1]	-1.759	-1.769	-1.728	-1.723
	(1.074)	(1.137)	(1.085)	(1.195)
RiceAdvice + Subsidy [T2]	-2.476	-2.347	-2.731	-2.614
	(1.837)	(1.740)	(1.856)	(1.731)
Combined treatment [T]	-4.235*	-4.117*	-4.459*	-4.338*
	(2.453)	(2.427)	(2.444)	(2.429)
Difference between treatments [T2-T1]	-0.717	-0.578	-1.003	-0.891
	(1.744)	(1.660)	(1.808)	(1.716)
Mean dependent variable in control	33.62			
Observations	623	623	592	592
R-squared	0.076	0.093	0.082	0.103
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes
Panel C: 2017 only				
RiceAdvice [T1]	-0.076	-0.108	0.173	0.216
	(0.901)	(0.898)	(0.951)	(0.976)
RiceAdvice + Subsidy [T2]	1.682	1.723	1.534	1.600
	(2.300)	(2.228)	(2.255)	(2.198)
Combined treatment [T]	1.606	1.615	1.707	1.816
	(2.662)	(2.602)	(2.690)	(2.661)
Difference between treatments [T2-T1]	1.758	1.831	1.360	1.383
	(2.263)	(2.185)	(2.178)	(2.118)
Mean dependent variable in control	32.45			
Observations	603	603	576	576
R-squared	0.207	0.219	0.213	0.225
LGA FE	Yes	Yes	Yes	Yes
Household covariates	No	Yes	No	Yes

Note: For simplicity, coefficient estimates are only reported for the treatment effect. All regressions include LGA (strata) fixed effects. ANCOVA estimates include baseline timing of urea application and year dummies, where applicable. Covariates include household size, age of household head, number of days in agricultural training, and indicators for if the household head has formal education, if farming is the household's main activity, and if they have access to credit. Standard errors clustered at the village-level are in parentheses ($***\ p < 0.01$, $**\ p < 0.05$, $*\ p < 0.1$).