Social Network Targeting of Agricultural Technology: Adoption, Input Substitution and Yield Effects

Andrew Dillon
Michigan State University
dillona6@anr.msu.edu

Maria Porter
Michigan State University
mporter1@anr.msu.edu

Aissatou Ouedraogo
Tufts University
aoued06@gmail.com

This Version: May 23, 2018

Selected Paper prepared for presentation at the 2018 Agricultural & Applied Economics Association Annual Meeting, Washington, D.C., August 5-August 7

Social Network Targeting of Agricultural Technology: Adoption, Input Substitution and Yield Effects[1]

Andrew Dillon
Michigan State University

Maria Porter
Michigan State University

Aissatou Ouedraogo
Tufts University

This Version: May 23, 2018

PRELIMINARY DRAFT

Abstract

Among subsistence farmers in low-income countries, the potential gains from agricultural technology adoption depend crucially on who adopts new technologies, how households internally reallocate resources amongst members, and whether yield gains are realized as a result of this reallocation. In the West African Sahel, few subsistence farmers adopt new seed varieties or fertilizer when planting sorghum, the main food staple and most widely cultivated dryland crop among these farmers. Farmers in the West African Sahel face a myriad of constraints in terms of available inputs and resources, which in our study, prove to be critical complements to production inputs such as fertilizer and improved seed. To address some of these constraints, we conducted a Randomized Control Trial in which we introduced sorghum farmers in Burkina Faso to a technology for applying small amounts of fertilizer at the time of planting, and provided a randomly assigned subset of farmers with free kits comprised of mineral fertilizer and improved seed. To identify effective methods for diffusing adoption of this technique, we randomized villages to either two forms of social network targeting of influential villagers, or to a more common non-targeted approach where farmers received free kits by random assignment. We find that the targeted approach is much more effective in encouraging broader adoption in the long-term. In terms of identifying potential gains to adoption, we find that smaller landholders who were already using mineral fertilizer saw considerable yield responses from our encouragement design.

Keywords: agricultural technology adoption, social networks, household decision-making
JEL codes: Q12, Q16, O12, O13, J00

[1] Acknowledgements: We are grateful to Michael Carter, Melinda Smale, Nicolo Tomaselli, Adama Traore and Chris Udry for very helpful comments and suggestions. We are very grateful to Estelle Plat for her excellent work as our IPA Research Associate, and to Ashesh Prasann for his research assistance. We also acknowledge financial support from the Bill and Melinda Gates Foundation and BASIS Feed the Future Innovation Lab for Assets and Markets Access. The project was reviewed for ethical approval by Michigan State University (xxx).

1. Introduction

Agricultural productivity growth is predicated on technology adoption. But the distribution of gains from technology adoption depends on who adopts new technologies, how households internally reallocate resources across plots, and whether yield gains are realized as a result of this reallocation. For example, in studying the effects of hybrid soybean seed adoption in Brazil, Bustos et al. (2016) found that initial household endowments such as weather and soil characteristics had heterogeneous effects on agricultural productivity. A large literature on technology adoption has focused on relaxing constraints to better understand which constraints[2] or market structures reduce agricultural investment (Duflo et al. 2008).

This paper contributes to this literature by first estimating the effects of social network targeting on adoption decisions, and secondly, estimating heterogeneity in household responses to such decisions.

In adopting a new agricultural technology, households respond differently in terms of related decisions on their use of other inputs. Such differences rely heavily on the additional constraints households face, such as the availability of complementary inputs to the new technology. As a result, the effectiveness of adopting a new agricultural technology in terms of increasing yields depends crucially on the availability of such complementary inputs.

This study provides a useful example for analyzing such dynamics, as we implement an encouragement design to promote the use of a labor-intensive fertilizer application technique

[2] For example, on credit and risk, see Karlan et al. (2014); on risk, see Cole et al. (2013); on social learning frictions within networks, see Beaman et al. (2015).

called microdosing. This technique is utilized at the planting stage of crop production, with farmers applying 2 to 6 grams of fertilizer (about a three-finger pinch) in or near the seed hole.[3] Microdosing contrasts with more commonly used broadcast methods of fertilizer application, where fertilizer is applied as top dressing from 3 to 4 weeks after the seeded crop begins to emerge. We provided sorghum farmers with a starter kit comprised of fertilizer and seed for a half hectare plot. We also provided training on the technique. As microdosing may often require more time at the planting stage, and as the amount of fertilizer was relatively small, we expect that more resource constrained farmers, with less available family labor, would see less gains to adopting the fertilizer application technique.

To understand how such adoption decisions may be influenced by social network targeting, we assigned each of the participating villages at random to one of three distribution strategies and a control group:

(1) Starter kits were randomly distributed to 15 sorghum farmers in the village (Group A). This group provides an estimate of the intent to treat effect of random distribution. It also provides a subsample of households for which we will estimate input substitution and yield effects in the second phase of analysis.

(2) We assigned 15 starter kits to farmers in each of these villages on the basis of their network size or degree. Farmers with the 15 largest network sizes in their villages received the starter kits in this second treatment group (Group B).

(3) In the third treatment group (Group C), we assigned 15 starter kits to farmers in each village on the basis of their network influence within the village as measured by eigenvector centrality (Jackson 2008).

[3] The amount of fertilizer applied is equivalent to about 20 to 60 kg of fertilizer per hectare.

(4) Finally, we also retained a control group of villages (Group D), where no training or kits were provided. These villages provide a counterfactual to the villages in (1) above.

We find that social network targeting was much more effective in encouraging sustained adoption of microdosing in the long-term compared to random distribution of starter kits. Whereas adoption rates at the plot-level were fairly similar across all three treatment groups in the year immediately following the free distribution of kits, when households were surveyed again four years after the initial distribution of kits, adoption rates had increased in the villages where social network targeting took place, and there was no difference in adoption rates between Groups A and D (random assignment and control group villages).

We find that while our intervention was effective in encouraging sorghum plot managers to use greater amounts of fertilizer, treatment effects varied both by plot size and by whether or not the farmer had already been applying mineral fertilizer to sorghum plots. Treatment intensity varied with plot size, because all households received the same amount of fertilizer, and households with smaller plots (below the median level) were more affected by being assigned treatment compared to those with larger landholdings. In terms of impacting fertilizer use, treatment effects were greatest for, and were maintained in the long-term, by small landholders who had not previously used mineral fertilizer. However, yield response was only notable in the short-term among the small landholders who had been using mineral fertilizer previously. In all, these results suggest that farmers who faced severe resource constraints in potentially complementary inputs to fertilizer and improved seed were limited in the potential gains they could derive from our free distribution of these inputs.

In the following section, we outline the experimental design. Section 3 describes the household demographics and production characteristics and balancing tests. Section 4 outlines the estimation methods. Section 5 summarizes the results and findings, and is followed by a discussion and conclusion.

2. Experimental Design

The focus of our study is to address the constraints to agricultural technology adoption among severely resource-constrained farmers who would potentially benefit from adopting microdosing, a technique that has proven to increase sorghum yields considerably in experimental plots (Ministry of Agriculture, Burkina Faso 2010). We therefore selected provinces in Burkina Faso where there was high prevalence of sorghum cultivation and little cultivation of cash crops. In doing so, we chose to conduct our study in Bam and Sanmatenga provinces from the Center-North region of the country, and Passoré province from the North region. . Like many rural households in the West African Sahel, households in these three provinces heavily rely on dryland sorghum cultivation as their main food staple. Although sorghum is a major source of livelihood for many households in Burkina Faso, national level average sorghum yields are estimated at 0.8 tons per hectare, despite the potential to attain over 2 tons per hectare (Ministry of Agriculture, Burkina Faso 2010).

A. *Household and Social Network Census*

As a main focus of our study is to understand the relative effectiveness of social network targeting in encouraging adoption, we collected social network census data to carry out such targeting. To do so, we restricted our sample of villages to those where such data collection would be feasible; that is, to villages that were not too large for surveying all households. We therefore restricted our study design to villages with a population size between 190 and 800, and where there were between 70 and 120 households living in the village in total. In the three selected provinces, this criteria left us with 925 villages in total.

We then randomly selected 80 villages for Group A, and 20 villages for each of the remaining groups outlined above (Groups B, C, and D). .

In each of 80 villages in Group A, approximately 15 households were randomly selected to receive free micro-packets of certified sorghum seed and fertilizer. We also provided training in microdosing application to all households in treated villages (Groups A, B, C).[4]

Within each village, all household heads and their members were fully enumerated in an initial visit. The village enumeration included questions about plot information, sorghum production, and adoption of improved seeds and fertilizer. To measure adoption decisions, we conducted a similar exercise in two follow-up waves of data collection: one year immediately following the

[4] The certification of the seed and training in microdosing application technique to all treated households were implemented by INERA, the public agricultural research institution in the country. The supply of the microdosing packets was performed by AGRODIA, in collaboration with INERA's trainers.

distribution of kits and training, and four years following this distribution and training. Attrition rates were relatively low in both follow-up rounds.

The original village enumeration and follow-up surveys were conducted in all 140 villages. However, a social network census was conducted for 100 of these villages. As social network information was required for targeting influential individuals in villages in Groups B and C, we collected social network census data in all 40 of these villages. We also did so in the control villages (in Group D). Due to budget limitations, we collected social network census data from 40 of the 80 villages in Group A. The relatively large number of villages in Group A enables us to identify the effects of random distribution of starter kits on adoption decisions, yields, other input substitution decisions, and potential secondary effects of the receiving kits. However, for analyzing adoption and diffusion of technology across households, we have data on 40 of these 80 villages for comparison to the 40 villages where kits were distributed to the most influential members.

From each household, a male and a female farmer who were considered to be the most knowledgeable in the household were asked to list members of the village who were in their social networks. These village members were identified on the basis of several criteria: those with whom they spoke frequently regarding agriculture, those with whom they had financial transactions, those who were their relatives or agricultural plot neighbors, and those who were members of the same organizations with which they were affiliated. All of these "social network links" were pooled together for each household..

We then used this data to calculate the total number of links for each household, as well as the household's eigenvector centrality, a measure of network influence (Jackson 2008). This information was used to identify recipients of starter kits in villages allocated to treatment Groups B and C.

B. *Household Survey*

To estimate input substitution and yield effects, a more detailed household survey of production activities was required. In each of the 80 villages in treatment Group A and in each of the 20 villages in control Group D, we randomly sampled approximately 30 households growing sorghum. Of these surveyed households, 15 households were randomly selected to received microdosing starter kits, and the remaining 15 surveyed households served as a counterfactual to our treated households. For all 30 households in each of the 100 villages, we conducted baseline and follow-up household surveys, where we collected detailed production and socio-economic information. Thus, detailed household surveys were conducted for 2400 households in all 80 villages in treatment Group A, and for 600 households in all 20 villages in control Group D. We conducted a second wave of follow-up households surveys in 40 of the 80 originally surveyed Group A villages (those where the social network census had been conducted), and all Group D villages. Attrition rates were low in both follow-up periods.

3. Household demographics and production characteristics

Tables 1 and 2 present summary statistics of pre-intervention household demographic and production characteristics for both the pooled sample (columns 1 and 2) and by treatment assignment (columns 3 through 6). In both Tables 1 and 2, columns 3 and 4 summarize statistics for untreated households in both treatment and control villages, while columns 5 and 6 summarize statistics for households assigned to treatment. To ensure that both demographic and production characteristics are balanced across treated and untreated households, we tested the mean equivalence across the two groups of households, with standard errors clustered at the village level; p-values for these means tests are reported in the last columns of Tables 1 and 2 and indicate that the randomization produced balanced assignment and control groups of households..

The statistics for the pooled sample in Table 1 show that on average, household heads are approximately 49 year old and only about 12% of the households are female-headed households. While the average number of adults in the household (those ages 15 years or older) is about 5.5, the average number of female members tends to be slightly higher than that of male members. One factor that might explain this is the high prevalence of polygamy, as more than 40% of the sample are polygamous households.

Household-level production characteristics are presented in columns 1 and 2 of Table 2. We report statistics for the household crop choice portfolio, land area allocated to crop cultivation, labor allocation by specific agricultural activity, fertilizer use, and crop yield performance. Note that we do not report statistics for sorghum as crop choice, because all households are sorghum growers, which was the primary criterion used in the sample selection. The statistics show that millet is the most significant crop other than sorghum, as more than 42% of households reported having at least one plot where millet is the main crop. The corresponding statistics for peanut and maize are about 12 and 10%.

The pattern in crop choice is also reflected in the total land devoted to the cultivation of these crops. While on average total land holding is about 4.6 hectares, the area of land devoted to sorghum production (as main crop) is about 3.5 hectares, resulting in up to 75% of the household total land holding being allocated to sorghum alone. These statistics are strongly reflective of the fact that sorghum constitutes the main staple food crop in the study area.

Descriptive statistics for labor allocation, expressed in person-days per hectare, are also presented in Table 2 for six agricultural activities ranging from land preparation (i.e., plowing) to harvest. On average, total labor allocation at the household level is about 86 person-days per hectare. Weeding tasks alone make up to 45% of this total labor allocation, making it by far the most labor intensive agricultural activity.

The last block of summary statistics in Table 2 summarize average crop yield performance at the household level and average rates of mineral fertilizer application. These statistics are unconditional on crop cultivation, which explains the very low level of yield for non-sorghum crops, especially, for rice, peanut and niebe, given the low proportion of households growing these crops. While all households in the sample cultivate sorghum, sorghum yield is only about 600 kilograms per hectare. This figure is significantly low compared to the reported national average yields for cereal crops, which is approximately 1,200 kilograms per hectare (World Bank, 2014).[5] Given that sorghum is the main staple food for these households, these statistics highlight the need for increased sorghum productivity in these regions.

---

[5] http://databank.worldbank.org/data/reports.aspx?source=2&series=AG.CON.FERT.ZS&country=

The proportion of households who applied mineral fertilizer (in the cultivation of any type of crop) before the program intervention was about 44%. Fertilizer application intensity, unconditional on fertilizer use, was about 17 kilograms per hectare. This figure, which is strikingly low, is similar to national level statistics reported by the World Bank.[6] Even after restricting the sample to only households who used fertilizer, the amount applied per hectare is only approximately 25 kilograms per hectare, (figure not reported here) while the standard recommended application amount varies between 100 and 200 kilograms per hectare.

4. Estimation Methods

We outline the methods used to estimate the intent to treat (ITT) effects of receiving microdosing kits and related training on microdosing adoption, input substitution and yield. In a second set of regression specifications, we also estimate the treatment on the treated effects on yield among those sorghum farmers who adopted microdosing.

A. Microdosing Adoption

We estimate adoption dynamics using two units of analysis, the household and the plot. At the household level, we estimate the intent to treat effect of village level social network targeting treatments based on random allocation, degree targeting or eigenvector centrality targeting as follows:

$$y_{hj} = \alpha + \beta_1\, Tr_j + \beta_2\, Td_j + \beta_3\, Tec_{hj} + \gamma X + \epsilon_{pghj} \qquad (1)$$

In this regression specification, y_{hj} is a microdosing knowledge or adoption indicator for household h in village j. If the household has correct knowledge about microdosing or has

[6] See the World Bank DataBank of World Development Indicators at: http://databank.worldbank.org/data/reports.aspx?source=2&series=AG.CON.FERT.ZS&country=

adopted microdosing on at least one of their plots, this variable takes the value of one and zero otherwise. The variable Tr_j is an indicator for whether the treated household resided in a village where kits and training were randomly assigned. The variable Td_j is an indicator for whether the treated households was selected according to degree while Tec_j is the indicator variable for an eigenvector centrality targeted village. The excluded category for these indicator variables is the control group households. The additional controls (X) includes household sorghum land size as the magnitude of the treatment will differ depending on the household's initial endowment of land. Standard errors are clustered at the level of the village to reflect the clustered treatment design.

In our second adoption specification, we estimate the intent to treat effect at the plot level to better understand adoption heterogeneity, particularly by gender of the plot manager:

$$y_{pghj} = \alpha + \beta_1\, Tr_j + \beta_2\, Td_j + \beta_3\, Tec_{hj} + \beta_4\, Female_{pghj} + \beta_5\, (Tr_{hj} \times Female_{pghj}) + \beta_6\, (Td_{hj} \times Female_{pghj}) + \beta_7\, (Tec_{hj} \times Female_{pghj}) + \gamma X + \epsilon_{pghj} \quad (2)$$

In this specification, y_{pghj} is a microdosing knowledge or adoption indicator for plot p managed by a farmer of gender g in household h in village j. If the plot manager has correct knowledge about microdosing or has adopted microdosing on their plot, this variable takes the value of one and zero otherwise. As in Equation 1, the variable Tr_j is an indicator for whether the treated household resided in a village where kits and training were randomly assigned. The variable Td_j is an indicator for whether the treated households was selected according to degree while Tec_j is the indicator variable for an eigenvector centrality targeted village. The excluded category for these indicator variables is the control group households. Regression (2) also includes the interactions between the three treatment group indicator variables and the gender of the plot manager. The additional controls (X) includes household sorghum land size as the magnitude of the treatment will differ depending on the household's initial endowment of land. Standard errors are clustered at the level of the village to reflect the clustered treatment design.

B. Input Substitution and Yield Effects

To capture potential dis-adoption dynamics, regressions are estimated separately for each of the two follow-up periods in 2014 and 2017. Since treatment was allocated to managers of sorghum plots, and since the majority of those managing sorghum plots are men, the sample is restricted to men in the household who managed sorghum plots prior to treatment.

We first estimate OLS regressions on an indicator for whether household *i* in village *j* received a kit in 2014:

$$Y_{ij}^{k} = \beta Treatment_{ij} + \gamma X_{ij} + \mu_j + \varepsilon_{ij}^{k} \quad (3)$$

Plot-specific outcomes (indexed by *k*) include the following: kilograms of fertilizer applied; labor in number of person-days; and kilograms of sorghum yield. All outcomes are in natural logarithms, and in per hectare terms. Since treatment is a household-specific indicator, standard errors are clustered by household to address potential correlation across plots within the same household.

Labor is included among the outcome variables because microdosing is a relatively labor-intensive technology, requiring significant amounts of time during the planting period, when fertilizer needs to be applied to each individual seed hole. Since households with greater numbers of members would have more labor available, the number of person-days per hectare is divided by the number of household members.

We control for the number of plots on which the plot manager planted sorghum as the main crop prior to treatment (X_{ij}). This accounts for the fact that those with more sorghum plots at baseline enter into the regresssion estimates more frequently compared to those with fewer sorghum plots at baseline.

Finally, we include village fixed effects (μ_j) in all specifications in order to control for village-specific unobservable factors that may influence outcomes and would therefore bias estimated treatment effects. These may include village-specific agronomic or weather-related factors, as well as cultural-specific norms or customs that may influence sorghum-related production decisions and outcomes. By including village fixed effects, the comparison group of plot managers to those who were received kits are other plot managers in the same village who did not receive kits but had access to training in microdosing.

The experimental design was an encouragement design to promote the use of mineral fertilizer among two groups of plot managers: those who were not already using mineral fertilizer, as well as those who were using less effective fertilizer application methods such as broadcasting. Approximately 55% of male plot managers did not use mineral fertilizer prior to receiving kits. To examine potential differences in treatment effects between these extensive and intensive margins, we estimate additional, separate regressions among plot managers who did not use mineral fertilizer prior to treatment, and those who did do so.[7]

Intensity of treatment also varies considerably by plot size. This is because all treated households received the same amount of fertilizer, regardless of the amount of land where sorghum was planted prior to treatment. As a result, those with relatively small plots devoted to sorghum received more fertilizer per hectare compared to those with larger sorghum plots. These differences in treatment intensity can be manifested both in terms of use of fertilizer and labor inputs, as well as sorghum yields.

Such differences across varying plot sizes would imply that treatment effects, particularly for yield per hectare, would vary considerably across varying plot size. To estimate such heterogenous treatment effects in terms of plot size, we interact the household-specific treatment assignment indicator with an indicator for whether the total number of hectares the household devoted to sorghum prior to receiving a kit was above the median across all households in the sample (τ). Maintaining the same notation as previously, we therefore estimate the following:

$$Y_{ij}^{k} = \beta_1 (Treatment_{ij} \cdot 1\{Plot\ Size_{ij} < \tau\}) + \beta_2 Treatment_{ij}$$
$$+\beta_3 \cdot 1\{Plot\ Size_{ij} < \tau\} + \gamma X_{ij} + \mu_j + \varepsilon_{ij}^{k} \qquad (4)$$

As in the previously outlined OLS regressions, outcomes of interest vary at the plot level, and standard errors are clustered by household. Additional covariates and sample restrictions from above are also used here.

[7] Note that including this indicator as an additional covariate in the wider sample yields similar results to those presented here.

In regression (4), the treatment effect for households with below average plot size is determined by the sum of coefficients β_1 and β_2. In contrast, the treatment effect of receiving microdosing kits for those with above average plot size is given by β_2. Thus, the difference in treatment effects across plot sizes is encompassed in estimates of β_1.

In addition to estimating the ITT effects of receiving a microdosing kit, we also estimate the Treatment on the Treated (TOT) effects of using mineral fertilizer. To do so, we restrict the sample to those households who were assigned treatment, and we estimate equivalent regressions to (3), where the treatment indicator is replaced by an indicator for having used mineral fertilizer. One important difference with these regressions, is that the indicator for having used mineral fertilizer is specific to the plot. We therefore estimate the following model:

$$Y_{ij}^k = \beta_1 Fertilizer_{ij}^k \cdot 1\{Plot\ Size_{ij} < \tau\}) + \beta_2 Fertilizer_{ij}^k + \beta_3 \cdot 1\{Plot\ Size_{ij} < \tau\} + \gamma X_{ij} + \mu_j + \varepsilon_{ij}^k \quad (5)$$

In this regression, fertilizer use is interacted with the same below median plot size indicator as used previously. As before, this indicator is based on the total landholdings managed by plot manager *i* at baseline, where sorghum was the main crop. We continue using this indicator here for several reasons: to maintain comparable heterogeneous treatment effect estimates to those in the ITT estimates; we estimate heterogeneous effects by this baseline characteristic, which must therefore be specific to the plot manager as we have plot manager panel data (not plot-specific panel data); and treatment was allocated to the household or plot manager, as it is not possible to enforce microdosing or fertilizer use on particular plots.

Results

We first present the microdosing adoption results from estimating regressions (1) and (2) at the household and plot levels respectively, before turning to the intent to treat and treatment on the treated results for input substitution and yield.

A. Microdosing Adoption

The household level adoption intent to treat results described in regression (1) are presented in Table 3. Microdosing knowledge is relatively high among households who were targeted based on their social network characteristics. Households that were targeted with random assignment of starter kits and training in treatment Group A had 7.2 percentage points higher knowledge than the control households, whereas households targeted with either degree or eigenvector centrality treatments were 26 percentage points and 23.4 percentage points, respectively, higher.

Higher microdosing knowledge did yield slightly higher adoption rates among the social network targeted groups, but not until several years after the initial starter kits were distributed. In the agricultural season that the starter kits were distributed, the change in microdosing between the baseline and the 2014 agricultural season was 6.7 percentage points higher relative to the control in the randomly assigned (T1) households, while adoption rates were 7.3 and 9.1 percentage points higher in the degree and eigenvector centrality treatments. Though there are only small differences in the magnitude of the adoption rates between treatments, each treatment effect is significantly different from the other according to our hypothesis tests presented in Table 3.

Where there is higher divergence between treatment groups in adoption rates is when we consider the change in adoption between the first agricultural season in 2014 and the 2017 agricultural season. Microdosing adoption continues to increase over this period, but by only 1.6 percentage points in the randomly assigned treatment group one. Adoption rates in both the degree and eigenvector centrality rates increase by 26.9 percentage points and 29.2 percentage points respectively.

In Table 4, we estimate the microdosing knowledge and adoption intent to treat effects at the plot level, estimating the treatment effect heterogeneity with respect to female farmers. We estimate these effects with the full sample in columns 1-3 and a restricted subsample of sorghum plots in columns 4-6. The knowledge and adoption patterns are similar with respect to the household level results presented in Table 3, though the magnitude of the effects is different as the number of plots are much larger from the census data than the number of households in the household sample. The treatment effects for all three treatments is slightly smaller when we consider only the restricted subsample of sorghum plots suggesting that some microdosing knowledge may have spilled over to other types of crops cultivated.

Of particular interest with the plot level results is the estimation of gender heterogeneity at the plot level. Microdosing knowledge was lower among female farmers when considering the full sample of plots in treatment group one (3.5 percentage points lower than men) and treatment group three (5.7 percentage points lower than men). In the sorghum plot restricted subsample, women’s knowledge of microdosing in 8.4 percentage points lower in the eigenvector centrality group. Microdosing adoption differences among female farmers were found primarily in the first year after the starter kit distribution, but not in the three subsequent agricultural seasons. This indicates that the lower adoption rates among women persisted in the subsequent seasons with female farmers not increasing adoption relative to men. In the first season after the distribution of the starter kits, women in the randomly assigned starter kit group were 1 percentage point less likely than men to adopt microdosing, while in the degree treatment group they were 2.5 percentage points less likely to adopt. These adoption rate differences were also estimated in the sorghum plot subsample.

B. Input Substitution and Yield Effects

Table 5 summarizes results for estimating OLS regression (3). Here, we see very significant treatment effects on the amount of fertilizer used on plots managed by male sorghum plot managers in the first year after kits were distributed. In particular, receiving a microdosing kit increases fertilizer use by 107 to 160 percent. With fertilizer use post-treatment averaging around 10 kilograms, these coefficient estimates imply an average increase of roughly 11 to 16 kilograms per hectare-per capita.

[TABLE 5]

Coefficient estimates are similar across the three different specifications, including for the smaller subsample of plots where fertilizer was not used at baseline. Thus, in terms of fertilizer use, treatment effects occurred at both the extensive and intensive margins, encouraging household members to use mineral fertilizer where they had not done so previously, as well as encouraging sorghum farmers to use more mineral fertilizer than previously.

With such strong treatment effects on fertilizer use, it is perhaps somewhat surprising that estimated treatment effects on labor and yield are so low in magnitude and imprecisely estimated. Yet, this underestimation of treatment effects under the assumption of strictly linear treatment effects appears to be due to the fact that treatment effects are in fact non-linear in plot size, which is consistent with the fact that treatment intensity varied depending on households' baseline landholdings.

These results for the first year after treatment (2014) contrast sharply with those in 2017. Four years after kits were distributed, there was no identifiable impact on fertilizer use or the other outcomes.

The results summarized in Tables 5 and 6 demonstrate that treatment effects were maintained through the final round of surveying for a specific sub-group of plot managers, namely, those who managed below median-sized sorghum plots and had not used any mineral fertilizer on their sorghum plots prior to the intervention.

Turning first to fertilizer as the outcome variable, treatment effects are higher for plots below the median plot size compared to those above. ITT estimates imply an increase in fertilizer use of 180 percent for plots below the median, and 91 percent for plots above the median. These differences across plot size are driven primarily by plot managers who did not use fertilizer prior to our intervention. For this sub-sample of plot managers, ITT estimates imply an increase in fertilizer use of 220 percent for plots below the median, and 103 percent for plots above the median. With fertilizer averaging around 10 kilograms per hectare-per capita, these estimates imply an increase of 22 kilograms per hectare for the smallest plots managed, and an increase of 10 kilograms for the largest plots.

The results also indicate that the majority of treatment effects for smaller landholders were due to the changes in fertilizer use at the extensive margin. Among plot managers who had already been using fertilizer prior to our intervention, there were also significant increases in fertilizer use. But among sorghum farmers with below median plot size, treatment effects at the intensive margin were significantly lower, with a 98% average increase in fertilizer use. Among sorghum farmers with above median plot size, treatment effects at the intensive margin were somewhat higher, at around 133%.

[TABLE 5]

While sorghum farmers increased the amount of fertilizer used on their fields, these significant increases in inputs did not necessarily translate into yield gains. One reason for this may be constraints in other inputs required for microdosing. As microdosing is a labor-intensive technique, increases in household labor would be required for the technique to be most effective. However, results in Table 5 suggest that treated households did not increase labor allocation. They may already be labor-constrained and have little available time to spare.

Despite the lack of response to treatment in terms of household labor, there was some considerable yield response due to treatment. Farmers with below median plots at baseline saw a 16% increase in yield if they received microdosing kits. In contrast, farmers with above median plots saw a 22% decrease in yield in response to receiving kits. These differences in yield response are most notable at the extensive margin, where smaller landholders who had already been using fertilizer previously saw a 24% increase in yield, while larger landholders saw a 26% decrease in yield (though the latter estimates are not statistically significant).

These heterogeneous treatment effects suggest that small landholders who are not severely constrained in terms of access to fertilizer and possibly other inputs were the ones to benefit from introduction of microdosing and the free distribution of inputs. However, these effects thus far have only been measured for the year immediately following the distribution of inputs.

In fact, these treatment effects were not maintained when farmers were surveyed again several years later. Results summarized in Table 6 show that there were no discernible treatment effects in terms of labor or yield for any sub-population of sorghum farmers affected. Since there were no labor effects in the first follow-up survey, it is not surprising that there also were none a couple of years later. The lack of yield response can be explained by the fact that there was no similar response to fertilizer use among small landholders who had already been using fertilizer prior to the intervention. Since this sub-population had been the main driver of the yield response in the first year, it is not surprising that there was no similar yield response in 2017.

[TABLE 6]

One reason for this may potential spillovers from those assigned free kits to other sorghum farmers in the village. Another reason for this is potential dis-adoption of microdosing three years after the initial intervention.

On a final note on the ITT estimates, it is notable that where we see a response to treatment four years after the intervention is among small landholders who had not previously used fertilizer. These resource-constrained sorghum farmers increased their fertilizer use by over 100% relative to similar farmers who had not been assigned kits. In contrast, larger landholders who had not used fertilizer in the baseline used 137% less fertilizer compared to similar farmers who had not been allocated treatment. Larger landholders may be more labor-constrained compared to smaller landholders, and the lack of yield response in the initial year may have turned them off from using fertilizer. While smaller landholders also did not see a discernible yield response in the first year, they may be less constrained to continue their efforts in using fertilizer. Clearly, there are many important dynamics at play in terms of adoption, dis-adoption, potential spillovers, and other resource constraints that may attenuate yield response to a labor-intensive fertilizer application technique.

The TOT estimates point towards evidence of spillovers to untreated households rather than dis-adoption as the likely explanation for the lack of discernible ITT effects three years following the intervention. Results are summarized in Table 7. In terms of fertilizer use, treatment effects were similarly high for both 2014 and 2017. TOT effects were also very similar across the intensive and extensive margins. In 2014, these effects were also similar across all plot sizes. There is some indication that treatment effects in 2017 were slightly higher for smaller landholders. These high and persistent treatment on the treated estimates provide no indication of dis-adoption over the years. Thus, the lack of discernible 2017 effects in the ITT estimates is likely due to adoption of microdosing and fertilizer by many of the untreated households.

[TABLE 7]

In terms of yield response, TOT estimates do indicate considerable heterogeneity, which again points towards further evidence of the importance of other resource constraints that have not been addressed by our intervention. In 2014, yield response was quite significant for larger landholders who had previously already been using fertilizer. Treatment estimates for other sub-samples of sorghum plot managers are negative. These results indicate that the least resource-constrained farmers were the only ones to benefit from receiving additional free inputs.

However, four years following the distribution of free kits, smaller landholders also benefited significantly from receiving such kits. While yield response was by far the highest among the least constrained farmers, small landholders also saw positive yield responses in 2017, with higher yield response at the intensive as opposed to extensive margin.

Conclusion

In sum, the results outlined above indicate that the most effective targeting policy for broadly diffusing a new agricultural technology among subsistence farmers in Burkina Faso would be to provide free inputs and training to the most influential members of the village. This may be because farmers who are socially connected to many others in the village would both encourage their friends to try out a new technology, and they may be more likely to see higher yield gains to doing so because they would be less resource-constrained compared to other village members; the lack of resource constraints would in terms of raising yields would be to target those who already use fertilizer, and therefore are not as severely resource-constrained as other farmers.

However, small landholders more broadly also stand to benefit greatly from encouragement designs such as ours. That small landholders who had already been using fertilizer saw yield gains following the free distribution of fertilizer and seed kits indicates that there is significant potential for raising yields for these more vulnerable farmers.

But in addition, it is critical to also address the other production constraints facing farmers. Farmers in the West African Sahel face a myriad of constraints in terms of available inputs and resources, which prove to be critical complements to inputs such as fertilizer and improved seed.

References

Aune, J.B. and Coulibaly, A. (2015). Microdosing of Mineral Fertilizer and Conservation Agriculture for Sustainable Agricultural Intensification in Sub-Saharan Africa. In Sustainable Intensification to Advance Food Security and Enhance Climate Resilience in Africa (pp. 223-234). Springer International Publishing.

Beaman, L., BenYishay, A., Magruder, J., and A.M. Mobarak. 2015. “Can Network Theory-based Targeting Increase Technology Adoption?” Northwester University working paper. Available at: http://faculty.wcas.northwestern.edu/~lab823/MNW_june15.pdf .

Bustos, P., Caprettini, B., and J. Ponticelli. 2016. “Agricultural Productivity and Structural Transformation: Evidence from Brazil.” American Economic Review 106(6): 1320-1365.

Cole, S., Gine, X., Tobacman, J., Topalova, P., Townsend, R., and J. Vickery. 2013. “Barriers to Household Risk Management: Evidence from India.” American Economic Journal: Applied Economics 5(1): 104-35.

Duflo, E., Kremer, M., and J. Robinson. 2008. “How High Are Rates of Return to Fertilizer? Evidence from Field Experiments in Kenya.” American Economic Review 98(2): 482-88.

ICRISAT, Fertilizer Microdosing, January 2009.

ICRISAT, Fertilizer Microdosing, July 2012.

Jackson, M.O. 2008. Social and Economic Networks. Princeton, NJ: Princeton University Press.

Karlan, D., Osei, R., Osei-Akoto, I., and C. Udry. 2014. “Agricultural Decisions after Relaxing Credit and Risk Constraints.” Quarterly Journal of Economics 129(1): 597-652.

Ministere de l’Agriculture, de l’Hydraulique et des Resources Halieutiques. 2010. Diagnostic du secteur Semencier du Burkina Faso. Document final. Ouagadougou, Burkina Faso, November.

Table 1. Pre-intervention household demographic characteristics

VARIABLES	All households (N=2,376) Mean	SD	Untreated households (N=1,435) Mean	SD	Treated households (N=941) Mean	SD	p-Value
Head age	48.75	15.28	48.92	15.34	48.49	15.19	0.54
Female head of household	0.12	0.32	0.13	0.33	0.10	0.30	0.05
Number of adult males	2.32	1.70	2.32	1.70	2.32	1.71	0.99
Number of adult females	2.79	1.90	2.82	1.84	2.74	1.98	0.40
Number of all adults	5.47	3.18	5.51	3.11	5.40	3.30	0.48

Table 2. Pre-intervention household production characteristics

	All households (N=2,376)		Untreated households (N=1435)		Treated households (N=941)		
VARIABLES	Mean	SD	Mean	SD	Mean	SD	P-Value
Crop choice, plot size, and land holding							
HH has plot(s) where main crop is millet	0.42	0.49	0.42	0.49	0.41	0.49	0.62
HH has plot(s) where main crop is maize	0.10	0.30	0.11	0.31	0.10	0.30	0.50
HH has plot(s) where main crop is rice	0.01	0.12	0.01	0.12	0.01	0.12	0.96
HH has plot(s) where main crop is peanut	0.12	0.32	0.10	0.30	0.13	0.34	0.02
HH has plot(s) where main crop is niebe	0.07	0.25	0.07	0.25	0.07	0.26	0.54
HH has plot(s) where main crop is other crop type	0.00	0.05	0.00	0.05	0.00	0.05	0.99
Number of plots	3.35	2.10	3.35	2.08	3.37	2.12	0.85
Number of sorghum plots	2.38	1.72	2.39	1.78	2.36	1.64	0.72
Total land holding	4.61	4.05	4.61	4.09	4.61	3.99	0.99
Total sorghum land	3.48	3.37	3.44	3.18	3.55	3.64	0.46
Total millet land	0.87	1.92	0.92	2.19	0.81	1.39	0.22
Total maize land	0.07	0.33	0.08	0.37	0.07	0.27	0.62
Total rice land	0.02	0.45	0.03	0.57	0.01	0.10	0.26
Total peanut land	0.09	0.36	0.08	0.34	0.11	0.38	0.08
Total niebe land	0.07	0.34	0.07	0.35	0.07	0.32	0.91
Total other crops land	0.00	0.03	0.00	0.04	0.00	0.02	0.40
Agricultural labor (person-days/ha)							
Total labor at plowing	10.10	10.26	10.22	10.35	9.92	10.12	0.55
Total labor at planting	13.95	11.31	14.07	11.37	13.78	11.22	0.58
Total labor at fertilizer application	2.62	4.89	2.51	4.85	2.79	4.95	0.24
Total labor at manure application	5.06	6.55	5.15	6.52	4.92	6.59	0.41
Total labor at weeding	38.37	36.25	39.22	36.12	37.08	36.43	0.18
Total labor at harvest	15.78	13.23	15.98	13.34	15.47	13.07	0.42
Total labor	85.97	62.31	87.24	62.71	84.04	61.67	0.22
Fertilizer application and crop yield in kg/ha							
Sorghum Yield	599.50	558.00	595.90	568.80	605.10	541.50	0.69
Millet yield	211.00	461.80	218.10	510.60	200.10	375.20	0.45
Maize yield	23.51	189.50	25.97	176.20	19.76	208.20	0.48
Rice yield	7.82	104.60	6.50	90.63	9.84	122.80	0.30
Peanut yield	34.53	167.10	30.18	165.60	41.17	169.20	0.05
Niebe yield	12.81	75.76	13.98	81.36	11.01	66.33	0.30
Fertilizer (1=if yes)	0.44	0.50	0.42	0.49	0.47	0.50	0.03
Fertilizer intensity (kg/ha)	8.29	17.29	7.63	16.96	9.28	17.75	0.06
Fertilizer price (FCFA/KG)	444.40	59.80	444.80	59.67	443.80	60.04	0.88

Table 3: Household Microdosing Adoption: Intent to Treat

	Microdosing Knowledge coef/se	Adopted Microdosing (2012-2014) coef/se	Adopted Microdosing (2014-2017) coef/se
T1: Random Assignment	0.072***	0.067***	0.016***
	(0.008)	(0.004)	(0.004)
T2: Degree Assignment	0.260***	0.073***	0.269***
	(0.013)	(0.010)	(0.015)
T3: EC Assignment	0.234***	0.091***	0.292***
	(0.014)	(0.011)	(0.016)
HH sorghum land size	0.002***	0.000***	0.001***
	(0.000)	(0.000)	(0.000)
Constant	0.535***	0.017***	0.027***
	(0.006)	(0.002)	(0.003)
Number of hh observations	13,865	14,466	14,466
pvalue_AB	0.000	0.000	0.000
pvalue_BC	0.000	0.000	0.000
pvalue_AC	0.000	0.000	0.000
pvalue_ABC	0.000	0.000	0.000

note: *** p<0.01, ** p<0.05, * p<0.1

Table 4: Plot-level Microdosing Adoption: Intent to Treat Effects

	Microdosing Knowledge	Adopted Microdosing (2012-2014)	Adopted Microdosing (2014-2017)	Microdosing Knowledge	Adopted Microdosing (2012-2014)	Adopted Microdosing (2014-2017)
	All Plots			*All Sorghum Plots*		
	coef/se	coef/se	coef/se	coef/se	coef/se	coef/se
T1: Random Assignment	0.059***	0.022***	0.001	0.054***	0.027***	0.003
	(0.011)	(0.004)	(0.006)	(0.012)	(0.004)	(0.006)
T2: Degree Assignment	0.057***	0.024***	0.134***	0.041**	0.026***	0.138***
	(0.016)	(0.007)	(0.014)	(0.018)	(0.008)	(0.015)
T3: EC Assignment	0.085***	0.019***	0.135***	0.063***	0.024***	0.132***
	(0.018)	(0.005)	(0.014)	(0.020)	(0.007)	(0.015)
Sex (1=Female)	-0.006	-0.016***	-0.024***	0.002	-0.016***	-0.024***
	(0.013)	(0.003)	(0.005)	(0.015)	(0.003)	(0.006)
T1 x Female	-0.035**	-0.010**	0.009	-0.033	-0.009*	0.000
	(0.017)	(0.004)	(0.007)	(0.021)	(0.005)	(0.008)
T2 x Female	0.019	-0.024***	0.014	0.021	-0.025***	0.031
	(0.027)	(0.007)	(0.018)	(0.034)	(0.008)	(0.028)
T3 x Female	-0.057*	-0.006	-0.000	-0.084**	-0.013	-0.032
	(0.031)	(0.008)	(0.018)	(0.040)	(0.009)	(0.022)
Land Size (ha)	0.012***	-0.001***	-0.001	0.012***	-0.002***	-0.002*
	(0.002)	(0.000)	(0.001)	(0.002)	(0.000)	(0.001)
Constant	0.466***	0.024***	0.063***	0.468***	0.028***	0.064***
	(0.010)	(0.003)	(0.006)	(0.011)	(0.004)	(0.006)
Number of plot observations	71,378	71,879	71,879	47,077	47,184	47,184
pvalue_AB	0.000	0.000	0.000	0.000	0.000	0.000
pvalue_BC	0.000	0.000	0.000	0.002	0.000	0.000
pvalue_AC	0.000	0.000	0.000	0.000	0.000	0.000
pvalue_ABC	0.000	0.000	0.000	0.000	0.000	0.000

Table 5. Estimated Intent to Treat Effects in 2014 (one year after kit allocation and training)

	Fertilizer (kg per hectare)			Labor (per-capita person-days per hectare)			Yield (kg per hectare)		
	All plots	No fertilizer used at baseline	Fertilizer used at baseline	All plots	No fertilizer used at baseline	Fertilizer used at baseline	All plots	No fertilizer used at baseline	Fertilizer used at baseline
Assigned kit x Below median baseline plot size	0.89	1.183	-0.347	0.176	0.036	0.194	0.387**	0.177	0.502**
	(0.544)	(0.720)	(0.852)	(0.119)	(0.174)	(0.176)	(0.151)	(0.209)	(0.230)
Assigned kit	0.909**	1.026*	1.326*	-0.118	0.013	-0.194	-0.226*	-0.17	-0.262
	(0.442)	(0.589)	(0.735)	(0.097)	(0.135)	(0.153)	(0.119)	(0.160)	(0.181)
Below median baseline plot size	-0.804**	-0.536	-0.536	0.033	0.023	0.075	-0.025	-0.014	0.059
	(0.353)	(0.461)	(0.564)	(0.075)	(0.101)	(0.115)	(0.088)	(0.122)	(0.141)
Num. Sorghum Plots at Baseline	-0.156	-0.159	-0.359	0.041	0.029	0.104***	-0.024	0.016	-0.024
	(0.132)	(0.177)	(0.231)	(0.025)	(0.037)	(0.038)	(0.031)	(0.041)	(0.051)
Observations	892	488	404	892	488	404	892	488	404
R-squared	0.2143	0.2724	0.2169	0.1705	0.2236	0.2921	0.1254	0.1569	0.2599
Treatment effect below median plot size	1.799***	2.209***	0.979*	0.058	0.049	0.000	0.161**	0.007	0.240*
Treatment effect above median plot size	0.909**	1.026*	1.326*	-0.118	0.013	-0.194	-0.226*	-0.170	-0.262

*** p<0.01, ** p<0.05, * p<0.1. Robust standard errors in parentheses, clustered by household. Village fixed effects are also included in all regressions. Sample is restricted to all male plot managers who managed a sorghum plot prior to treatment and who were also surveyed in both followup periods. Outcome variables are plot-specific, kit assignment and baseline plot size are household-specific, and the number of sorghum plots and fertilizer use at baseline are specific to plot managers. Fertilizer use at baseline refers to baseline sorghum plots only. All outcomes are in natural logarithms (with 0.01 replacing zero's).

Table 6. Estimated Intent to Treat Effects in 2017 (four years after kit allocation and training)

	Fertilizer (kg per hectare)			Labor (per-capita person-days per hectare)			Yield (kg per hectare)		
	All plots	No fertilizer used at baseline	Fertilizer used at baseline	All plots	No fertilizer used at baseline	Fertilizer used at baseline	All plots	No fertilizer used at baseline	Fertilizer used at baseline
Assigned kit x Below median baseline plot size	1.332**	2.428***	-0.393	0.067	-0.017	0.141	0.148	-0.014	0.227
	(0.529)	(0.767)	(0.770)	(0.155)	(0.216)	(0.245)	(0.224)	(0.242)	(0.384)
Assigned kit	-0.966**	-1.374**	-0.329	-0.175	-0.263*	-0.146	-0.29	-0.093	-0.383
	(0.416)	(0.571)	(0.639)	(0.124)	(0.159)	(0.213)	(0.191)	(0.186)	(0.351)
Below median baseline plot size	-0.185	-0.296	0.694	0.068	0.068	0.141	0.195	0.217	0.292
	(0.324)	(0.444)	(0.490)	(0.102)	(0.162)	(0.146)	(0.124)	(0.152)	(0.234)
Num. Sorghum Plots at Baseline	-0.069	-0.072	-0.102	-0.053	-0.029	-0.053	-0.071	-0.04	-0.058
	(0.119)	(0.162)	(0.210)	(0.037)	(0.055)	(0.058)	(0.044)	(0.058)	(0.074)
Observations	892	488	404	892	488	404	892	488	404
R-squared	0.2491	0.2672	0.2417	0.1598	0.2222	0.2459	0.1392	0.2098	0.2011
Treatment effect below median plot size	0.366**	1.054***	-0.722	-0.108	-0.280	-0.005	-0.142	-0.107	-0.156
Treatment effect above median plot size	-0.966**	-1.374**	-0.329	-0.175	-0.263*	-0.146	-0.290	-0.093	-0.383

*** p<0.01, ** p<0.05, * p<0.1. Robust standard errors in parentheses, clustered by household. Village fixed effects are also included in all regressions. Sample is restricted to all male plot managers who managed a sorghum plot prior to treatment and who were also surveyed in both followup periods. Outcome variables are plot-specific, kit assignment and baseline plot size are household-specific, and the number of sorghum plots and fertilizer use at baseline are specific to plot managers. Fertilizer use at baseline refers to baseline sorghum plots only. All outcomes are in natural logarithms (with 0.01 replacing zero's).

Table 7. Estimated Treatment on the Treated Effects in 2014 and 2017

	Fertilizer (kg per hectare)			Yield (kg per hectare)		
	All plots	No fertilizer used at baseline	Fertilizer used at baseline	All plots	No fertilizer used at baseline	Fertilizer used at baseline
2014 (One Year After Kit Allocation and Training)						
Fertilizer Used x Below median baseline plot size	0.279	0.28	0.318	-0.419*	0.015	-1.041**
	(0.202)	(0.388)	(0.309)	(0.225)	(0.314)	(0.430)
Fertilizer Used	7.076***	6.954***	7.112***	0.307*	-0.151	1.035***
	(0.175)	(0.322)	(0.236)	(0.175)	(0.195)	(0.361)
Below median baseline plot size	0.119	0.175	0.127	0.655***	0.173	1.310***
	(0.109)	(0.165)	(0.209)	(0.205)	(0.231)	(0.380)
Num. Sorghum Plots at Baseline	0.110**	0.122	0.138*	0.009	0.190***	0.007
	(0.042)	(0.083)	(0.078)	(0.053)	(0.069)	(0.081)
Observations	305	152	153	305	152	153
R-squared	0.9545	0.9572	0.9651	0.1931	0.3389	0.3821
Treatment effect below median plot size	7.355***	7.234***	7.430***	-0.112	-0.136	-0.006**
Treatment effect above median plot size	7.076***	6.954***	7.112***	0.307*	-0.151	1.035***
2017 (Four Years After Kit Allocation and Training)						
Fertilizer Used x Below median baseline plot size	0.765***	0.556	0.757*	-0.114	0.365	-1.495
	(0.242)	(0.437)	(0.455)	(0.419)	(0.429)	(1.534)
Fertilizer Used	6.503***	6.560***	6.641***	0.749**	0.301	2.421*
	(0.214)	(0.395)	(0.390)	(0.363)	(0.344)	(1.450)
Below median baseline plot size	-0.009	0.112	-0.04	0.395	-0.018	1.678
	(0.110)	(0.175)	(0.346)	(0.382)	(0.308)	(1.436)
Num. Sorghum Plots at Baseline	-0.04	0.175**	-0.142	-0.084	0.025	-0.212*
	(0.056)	(0.078)	(0.112)	(0.057)	(0.084)	(0.117)
Observations	305	152	153	305	152	153
R-squared	0.9354	0.9603	0.9198	0.2338	0.3344	0.3512
Treatment effect below median plot size	7.268***	7.116***	7.398***	0.635***	0.666**	0.926**
Treatment effect above median plot size	6.503***	6.560***	6.641***	0.749**	0.301	2.421*

*** p<0.01, ** p<0.05, * p<0.1. Robust standard errors in parentheses, clustered by household. Village fixed effects are also included in all regressions. Sample is restricted to all male plot managers who managed a sorghum plot prior to treatment, were assigned a kit or treatment, and who were also surveyed in both followup periods. Outcome variables are plot-specific, kit assignment and baseline plot size are household-specific, and the number of sorghum plots and fertilizer use at baseline are specific to plot managers. Fertilizer use at baseline refers to baseline sorghum plots only. Fertilizer use as a regressor is plot-specific. All outcomes are in natural logarithms (with 0.01 replacing zero's).