Are Matching Funds for Smallholder Irrigation Money Well Spent?

Conner Mullally^{§,†}

Shourish Chakravarty[‡]

This version: March 2018

Abstract

Groundwater irrigation can dramatically affect agricultural production and productivity. Despite its potential as an agricultural development tool, little credible evidence exists for the impacts of groundwater development on smallholder agriculture. We add to the evidence on the benefits of irrigation investments for small producers by evaluating the Rural Business Development (RBD) program of the Millennium Challenge Corporation in Nicaragua for small plantain producers. The RBD program offered matching funds covering up to 30% of the cost of two years of inputs, extension services, and diesel-powered micro-sprinkler irrigation for individual farms. In order to estimate the average impact of the RBD program on its beneficiaries, we combine model selection via the LASSO with doubly robust treatment effect estimation as applied to two years of panel data for 146 producers. We find that the program had substantial impacts on plantain revenue and production, while achieving nearly complete irrigation coverage of plantain land among beneficiaries. We find no discernible impact on household expenditure.

JEL Codes: O13, O22, Q12

Keywords: Agricultural development, impact evaluation, irrigation, machine learning

[§] Corresponding author, connerm@ufl.edu

[†] Assistant Professor, Food and Resource Economics Department, University of Florida. Conner Mullally gratefully acknowledges funding from the Millennium Challenge Corporation and the BASIS Research Program at the University of California, Davis.

[‡] Ph.D. Candidate, Food and Resource Economics Department, University of Florida

1. Introduction

While many interventions can positively affect agricultural production, the history of agricultural development shows that irrigation can be transformative. For example, Edwards and Smith (2017) estimate that expansion of irrigation accounted for 90% of the growth in the value of crop production in the Western United States after 1940. Fan, Gulati, and Thorat (2008) find that each rupee invested in irrigation returned an average of eight rupees in agricultural GDP in India throughout the 1960s and 1970s. Irrigation may be particularly beneficial for smallholder producers of perennial crops, as access to water can be constrain participation in modern agricultural value chains. Supermarkets and exporters require a constant and reliable supply of product (Barrett et al. 2012), and smallholder producers of perennial crops relying on rainfall may fail to meet this criterion. If small producers underinvest in irrigation because of difficulty accessing credit or other constraints, then promoting irrigation among small producers may improve welfare.

In this paper, we add to the evidence on the effectiveness of smallholder irrigation investments by measuring the impacts of the Rural Business Development Program (RBD) for small plantain producers in Nicaragua. The RBD program paid up to 30% of the cost of executing producer “business plans”. Business plans consisted of installing micro-sprinkler irrigation systems, purchasing inputs such as fertilizer, planting material, and pesticide, and obtaining extension advice on fertilizer use, integrated pest management, irrigation, as well as harvest and post-harvest management. In addition, the RBD program expanded capacity at a local collection center in order to address post-harvest bottlenecks.

Using a two-year panel data set of 95 RBD beneficiaries and 51 non-beneficiaries, we estimate the Average Treatment Effect on the Treated (ATT) for a series of outcome variables, where the ATT is identified by the Conditional Independence Assumption (CIA). In estimating the ATT, we follow a procedure detailed by Farrell (2015) that combines “doubly robust” treatment effect estimation with model selection using the LASSO algorithm. We use the LASSO to select from a set of variables that are potentially correlated with our outcomes of interest as well as RBD participation under a unitary agricultural household model. Doubly robust estimation combines regression with inverse probability weighting, and identifies the ATT if either the propensity score model or the model for the conditional mean of the untreated potential outcome is correctly specified. By relying on the LASSO for variable selection, we minimize ad hoc modeling assumptions and select variables in a way that results in valid inference under standard conditions, in contrast to many other model selection procedures.

In year one, we find that the RBD program increased irrigated plantain area by 42%, reduced exposure to drought by 16%, and reduced fertilizer expenditure per unit of land by 33%. We find no clear evidence for or against program impacts on other outcomes. In year two, we find large and statistically significant impacts on sales and production indicators, as well as irrigated area. Larger and more precisely estimated impacts in year two were most likely a result of program design, as beneficiaries received irrigation late in the first year of participation. Year one impacts would have been caused by input subsidies and extension. As a result of RBD participation, plantain revenue, production, revenue per unit of land, and yield grew by 249%, 130%, 110%, and 46% in year two, respectively. The

value of credit received by beneficiaries in year two fell by 60% as a result of RBD participation. The RBD program increased irrigated plantain area and area producing plantain among beneficiaries by 72% and 24%, respectively, in year two. As a result of the RBD program, nearly 93% of area producing plantain on beneficiary farms was irrigated, as compared to an estimated 66% in the absence of the program. Our results are robust to a large number of changes in model specification.

While agricultural indicators grew sharply as a result of RBD participation, we find no discernible impact on household expenditure. We consider several potential reasons for this finding, and cautiously conclude that the most likely explanation is a sharp increase in production costs. Administrative cost data provided by RBD program administrators show that variable costs of production increased from around \$800 per beneficiary to over \$4,000. Beneficiaries also contributed an additional \$2,500 in year one for irrigation equipment and land preparation. Whether beneficiaries will enjoy higher expenditures in future years is unclear. On one hand, the high cost of land preparation and irrigation equipment may indicate that beneficiaries traded short-term expenditure gains for long-term benefits. On the other hand, the sharp growth in variable costs may reflect modest increases in profits in future years.

In general, the available empirical evidence demonstrates that access to irrigation has positive effects on yield, production, and revenue, as well as household expenditure or income. Examples of studies finding positive effects of irrigation access include Huang et al. (2006), Del Carpio, Loayza, and Datar (2011), Dillon (2011a, 2011b), Kuwornu and Owusu (2012), Van den Berg and Ruben (2006), Hagos et al. (2012), and Adeoti et al.

(2009). Most existing studies measure impacts of surface water projects serving large numbers of farmers or the effects of on-farm irrigation technologies designed for larger farmers. Studies of the impact of treadle pump technology are exceptions (e.g. Adeoti et al. 2009). We address a gap in the agricultural development literature by evaluating a program that promoted smallholder adoption of motorized groundwater irrigation, a technology with far greater potential for production impacts than human-powered pumps.

In addition, our paper is one of the few to use a credible identification strategy in evaluating a matching funds program designed to promote agricultural development. Matching funds are a common development policy tool, both inside and outside of agriculture (Crespi and Maffioli 2014). As shown by recent reviews of the literature on matching funds programs in agriculture, the quality of the available evidence for their effectiveness is generally poor (Ton et al. 2013; Ton et al. 2015; Nankhuni and Paniagua 2013). In particular, evidence for the effectiveness of matching funds for investment in irrigation is virtually non-existent.

In what follows, we first present background on the development of plantain production on Nicaragua's Pacific Coast and the RBD program. We then describe the data set and present summary statistics for program beneficiaries and non-beneficiaries in our sample. Next, we discuss our identification strategy and model selection process before presenting our results and summarizing several robustness checks. We conclude with a discussion of our findings and possible directions for future research.

2. Background

Producers on the Pacific Coast of Nicaragua began farming plantain as a main crop beginning in the 1990s on land received through earlier agrarian reforms. Production was limited to small-scale operations where farmers sold output individually to intermediaries at the farm gate. Beginning in 2003 farmers began to organize themselves into cooperatives with the help of development projects from the EU and USAID. Greater organization is a step towards deeper involvement in modern agricultural value chains, as it can lower the transaction costs associated with dealing with large numbers of small producers (Barrett et al. 2012). However, to participate fully in modern agricultural value chains, small producers must meet the quantity and quality standards demanded by grocery stores, exporters, and processors.

The RBD program for plantain producers was introduced in 2007 in order to continue the agricultural development process described above. The plantain program was part of a larger development package for the Nicaraguan Pacific Coast co-financed by the Millennium Challenge Corporation (MCC), which is a development agency of the US government, and the Government of Nicaragua (GON). The Pacific Coast was selected by MCC and the GON for its economic potential, particularly in agriculture. Other components of the compact between MCC and GON included extensive road

rehabilitation, investment in several agricultural value chains, and a land titling program that was canceled following the 2007 Nicaraguan presidential election.^{1,2}

The RBD plantain program was administered by MCC's office in Nicaragua, known as the Millennium Challenge Account in Nicaragua (MCA-N), and Chemonics, a private firm specializing in the management of development projects. MCA-N and Chemonics promoted the RBD program in its initial stages through meetings with producers groups and cooperatives. Cooperative membership was not a prerequisite for RBD participation. Three cohorts entered the RBD plantain program from 2007 through 2009. The results presented in this paper include impacts of the RBD plantain program in 2010 and 2011 for the cohort entering in 2009, which consisted of 239 producers. RBD program services were co-financed by beneficiaries and the program itself. For each individual producer, the RBD program financed up to 30% of participation costs. For the program as a whole, the program paid 24% of all costs while facilitating access to the inputs, irrigation equipment, and other materials provided to beneficiaries.

Any farmer meeting a set of eligibility criteria could submit a "business plan" to the MCA-N office. The full set of criteria is provided in the supplementary appendix. Each business plan detailed how the producer would benefit from participation, how he or she

¹ The background information presented here draws from RBD program documents.

² See Carter, Toledo, and Tjernström (2016) for an evaluation of other components of the MCC-Nicaragua compact.

would finance participation, evidence that he or she had no access to services similar to those provided by MCA-N, and a detailed budget. MCA-N provided farmers with technical support in designing their business plans. While detailed data on applicant rejections are not available, conversations with MCA-N and Chemonics indicate that lack of financing and insecure property rights over land producing plantain were the most common reasons for rejection.

If a business plan was approved, the beneficiary received inputs (planting material, fertilizer, and agrochemicals) as well as extension services in each year of participation. Beneficiaries also received micro-sprinkler irrigation systems late in the first year of participation. Inputs and irrigation were meant to be sufficient for a single manzana³ of land. For the cohort studied in this paper, extension agents emphasized fertilization, integrated pest management, and management of the irrigation system.

3. Data

The data set used in our analysis consists of a two-year panel of 146 plantain producers, including 95 RBD participants and 51 non-participants serving as a control group.⁴ RBD

³ 1 manzana = 1.72 acres = 0.70 hectares

⁴ While 150 producers were interviewed for the initial round of data collection, three were not included in the follow-up round. One producer was dropped from our analysis for

participants included in the sample were randomly selected from a roster provided by MCA-N, whereas the comparison group was selected from a list of rejected applicants as well as through random sampling of non-participants with farms located in the same municipalities as RBD participants. All non-participants were screened to ensure that they satisfied the RBD eligibility criteria. Households were interviewed in the first quarter of 2010 and again in the second quarter of 2011.

The survey instrument was designed to collect data on outcomes likely to be affected by the RBD program, including household consumption (calculated as in Deaton and Zaidi (2002)), plantain production and revenue, irrigated area, area planted with plantain, area planted with other crops, value of credit received, self-reported exposure to plantain production shocks, and inputs applied to plantain (although input data were only collected in the first round of interviews). Area in maize and total credit were included in the set of outcomes to determine whether the program might have crowded-in or crowded-out other activities.

The survey instrument was also written to measure variables that might be correlated with outcomes of interest or participation in the RBD program under a unitary model of a risk-averse agricultural household making decisions in the presence of market failures. In order to avoid biasing estimated impacts by controlling for variables potentially

exceeding the maximum farm size allowed by RBD program rules. We use the balanced panel in our analysis.

affected by RBD participation (Angrist and Pischke 2009), the survey included recall questions on household membership, agricultural and non-agricultural wealth, ownership of non-agricultural businesses, credit access, experience with plantain, and land use. Demographic information captured in the survey and used in our analysis includes household size, the sex of the household head, and education levels of household members. If labor markets are imperfect, larger households may be better able to provide the labor needed for land preparation (expansion of plantain area was a goal of the RBD program) as well as maintaining additional plantain while harvesting increased output. The sex of the household head could be correlated with the ability of the household to access inputs complementary to those provided through RBD participation (Doss 2014). A common finding in the literature is that education is positively correlated with adoption of new technologies (Foster and Rosenzweig 2010), and therefore we might expect education to be correlated with RBD participation and effective use of irrigation.

Despite the fact that the RBD program helped finance irrigation investments and provided extension to assist with learning, wealth might still play a role in the ability to match funds received through the program or bear the risk associated with potentially unfamiliar irrigation technology (Foster and Rosenzweig 2010). Similarly, credit access may be key to financing the portion of RBD investments not covered by the program. Lastly, experience with plantain, land use, and ownership of non-agricultural businesses may indicate the degree of specialization in plantain production for a given household. We would expect households with greater specialization in plantain to have relatively high interest in the RBD program, and possibly have better performance as plantain farmers.

[TABLE 1 HERE]

Table 1 presents summary statistics for RBD participants and non-participants based on the first round of survey data. All variables except *Total unearned income* are based on recall questions. Since *Total unearned income* measures passive income sources such as retirement pensions and government transfers, it seems implausible that it would have been affected by program participation, but it could serve as a source of liquidity for plantain production. Therefore we include *Total unearned income* measured in the first survey round as a potential control variable.

As shown in the first four rows of the table, the demographic characteristics of the two subsamples are quite similar. Non-participants are about twice as likely to own a non-agricultural business (p -value = 0.019), while both groups belong to producer cooperatives in approximately the same proportion. The main plantain plots of farmers in the sample are located in two municipalities, and the summary statistics for *Farm located in municipality 1* suggest that there are no major differences in geographic location for farms in the two subsamples. RBD participants own about 26% more land than non-participants (p -value = 0.015), while number of plots and indices of other wealth sources (obtained through principal components analysis) are similar across the two subsamples.

Non-participants received about 28% more in loans than participants prior to the start of RBD for the cohort studied in this paper, but the difference in average credit value is small relative to the standard deviation in each subsample. RBD participants received about 43% more in unearned income, but the average amount is very small relative to the

value of credit received. Average experience with plantain, measured as the number of years out of the last ten that the household had plantain, is almost equal across the two subsamples. The Simpson index of crop diversity, calculated by squaring the proportion of planted area accounted for by each crop and then summing across crops, indicates that RBD participants and non-participants have crop portfolios that are about equally diversified. RBD participants dedicate 68% more area to maize (the dominant crop other than plantain) than non-participants, while non-participants have about 18% more land in plantain. Non-participants also have approximately 33% more irrigated area on their plantain plots at baseline.

In general, the differences in baseline characteristics among RBD participants and non-participants observed in table 1 are small, either in absolute terms or relative to the spread of the data. The exceptions to this pattern are *Proportion that owns a non-agricultural business* and *Land area*. As described in the next section, we assume that we can obtain unbiased estimates of RBD impacts by controlling for observed characteristics.

4. Empirical strategy

We begin by describing our identifying assumptions, before moving on to estimation, inference, and robustness checks.

4.1 Identification

Let $d_i = 1$ if producer i is an RBD participant and $d_i = 0$ otherwise. The treatment indicator, d_i , does not have a time subscript because participation status is fixed over time for each producer. Let y_{it}^d be the value of an outcome variable that would be observed for

producer i at time t if he or she had RBD participation status d . Let $\mathbf{x}_i$ be a vector of covariates and transformations of covariates. We assume that the following identifying assumptions hold.

$$E[y_{it}^0 \mid d_i, \mathbf{x}_i] = E[y_{it}^0 \mid \mathbf{x}_i] \quad (1)$$

Equation (1) is a version of the Conditional Independence Assumption (CIA), and it states that the untreated potential outcome is mean independent of treatment status. Mean independence is paired with an overlap condition on the conditional probability of treatment, i.e. the propensity score:

$$p(d_i = 1 \mid \mathbf{x}_i) < 1 \text{ for all } \mathbf{x}_i \in \mathbf{X}_1 \quad (2)$$

where $\mathbf{X}_1$ represents the set of all values of the vector of covariates observed among households participating in the RBD program. Equation (2) ensures that there are valid control households for all treated households. When combined with the Stable Unit Treatment Value Assumption (SUTVA), i.e. there are no program spillovers and all beneficiaries receive the same intervention, the conditions given in (1) and (2) are sufficient to identify the ATT of the RBD program.

The two channels through which SUTVA might be violated by the RBD program include general equilibrium price effects and information spillovers. The relatively small number of RBD beneficiaries suggests that the RBD program would not have had general equilibrium price effects. While we would expect irrigation investments to account for the vast majority of RBD program impacts, we cannot completely rule out farmer-to-farmer

transmission of knowledge gained from extension. To the extent that RBD extension had positive effects, it is possible that our ATT estimates are somewhat conservative.

4.2 Estimation

Taken together, the identifying assumptions listed above (which we will refer to as the CIA) ensure that unbiased and consistent estimation of the ATT is possible using regression, matching methods, inverse probability weighting, or combinations thereof (Imbens 2004). Although the CIA is silent on key issues of functional form, economic theory and contextual knowledge fill this gap by providing a guide as to which variables satisfy an exclusion restriction and therefore need not be included in data collection or estimation. In the present case, we have assumed that a unitary household model adequately describes the decision making process underlying RBD participation and each outcome of interest. This assumption carries with it a set of exclusion restrictions, e.g. intra-household bargaining strength can be excluded from all models without biasing estimation results, which enable us to proceed with estimation of the average impacts of the RBD program with the available data.

However, guidance from theory may be insufficient for determining the specifics of variable selection or model functional form. For example, even if theory presents a set of variables to hold fixed in estimation, it is difficult to know for certain whether one ought to control for higher order terms. As a result, model selection techniques are often informal, consisting of experimenting with a variety of specifications until discovering what “works” without any clearly specified criteria for what represents a satisfactory model. Post model-

selection inference is usually invalid if the effects of pretesting are not taken into account (Danilov and Magnus 2004).

The above facts motivate a recent literature applying machine-learning (ML) algorithms such as the LASSO to the estimation of average treatment effects (Belloni, Chernozhukov and Hansen 2014). On their own, the LASSO and related methods are no more than variable selection algorithms used to build models with excellent predictive properties. It is only when paired with an identification strategy that the LASSO and other ML algorithms can serve as powerful tools for treatment effect estimation. (Belloni, Chernozhukov and Hansen 2014).

The LASSO constrains the sum of the absolute values of the regression coefficients to be less than a “tuning parameter,” thereby forcing some coefficients to zero. The LASSO can be used even if the pool of covariates is larger than the sample size. Loosely speaking, as long as the true conditional mean function can be closely approximated by a model using fewer covariates than observations, the LASSO can serve as a highly accurate model selection algorithm. Furthermore, it is straightforward to construct accurate confidence intervals for causal effects of interest after performing model selection with the LASSO (Belloni, Chernozhukov and Hansen 2014; Farrell 2015).

We estimate and perform hypothesis testing on the ATT of the RBD program using a method developed by Farrell (2015). The method includes separate steps for model selection, estimation, and inference. In the model selection stage, the LASSO is applied to a logistic model of the propensity score using the entire sample, and then to a regression

model (logit for binary outcomes, linear for continuous outcomes) for each outcome of interest using only the control group. For the propensity score and each outcome, we obtain LASSO results under 100 different values of the tuning parameter. By iterating through successively smaller values of the tuning parameter, we increase model complexity from one parameter (i.e. the intercept) to $N + 1$ parameters, where N is the size of the estimation sample. The LASSO constraint ensures identification of the model with $N + 1$ parameters.

The covariates with non-zero coefficients under each value of the LASSO tuning parameter define alternative sets of covariates that could be used in ATT estimation. We compare the alternative sets of covariates using five-fold cross validation. The set of covariates that minimizes the sample average of a goodness of fit statistic generated through cross validation is used in ATT estimation. For the propensity score model, our goodness of fit measure is the deviance:

$$\frac{-2}{N} \sum_{i=1}^N [d_i \mathbf{x}_i' \hat{\beta} - \ln(1 + \exp(\mathbf{x}_i' \hat{\beta}))] \quad (3)$$

For each observation i , the estimated parameter vector $\hat{\beta}$ is obtained using the cross validation blocks that do not include i .

The cross validation procedure for our outcome variables proceeds in a similar fashion. The major differences are that the sample is restricted to households that did not participate in the RBD program and the outcome models are estimated by OLS when the dependent variable is continuous. When evaluating linear models, our goodness of fit statistic is mean squared prediction error for non-participant households:

$$\frac{1}{\sum_{i=1}^N (1 - d_i)} \sum_{i=1}^N (1 - d_i) (y_i - \hat{y}_i)^2 \quad (4)$$

For each observation i , the fitted value $\hat{y}_i$ is generated from a model estimated using the four cross validation folds that do not include i . For binary outcomes, the goodness of fit statistic is the deviance given in equation (3), calculated using non-participants.

The pool of covariates entered into the LASSO algorithm includes the variables shown in table 1 as well as several additional controls. First, we add additional controls for geography to complement our indicator for the municipality in which the household farm is located. During data collection, GPS coordinates and the altitude of the main household plantain plot were recorded. We use the GPS data to create an indicator of geographic location by splitting the latitude and longitude coordinates at their sample medians and placing each plot into the appropriate quadrant created by the split. Plots are only located in two of the resulting four quadrants, and we include a dummy variable for having a plot in one of the two quadrants in our pool of covariates, as well as the natural log of the altitude of the plot above sea level.

We also include two additional factors from our principal components analysis of household agricultural wealth variables, as well one additional factor for non-agricultural wealth. For agricultural as well as non-agricultural wealth, we select the number of factors by plotting the eigenvalue of each factor against the rank obtained when ordering the eigenvalues from largest to smallest, and then picking all factors above any obvious drop off in the size of the eigenvalues (Zhu and Ghodsi 2006). The eigenvalue plots are included

in the appendix. We limit the number of factors in order to save on the computational time associated with the robustness checks performed in our results and to avoid introducing additional noise into our estimates. Our main results are robust to marginal increases in the number of retained factors. Lastly, the pool of covariates used in estimation includes all squared non-binary variables and a full set of interactions for the variables in table 1 as well as those listed in the preceding paragraphs. In total, we have a pool of 292 covariates.

Once the final set of covariates is selected for the propensity score model as well as each outcome, the logit and regression models used to obtain the ATT are estimated without the LASSO constraint in order to avoid biasing model coefficients downwards. In other words, we use the “post-LASSO” estimator rather than the constrained estimator to construct our treatment effect estimator (Belloni, Chernozhukov and Hansen 2014; Farrell 2015; Hastie, Tibshirani, and Wainwright 2015). The estimated ATT for outcome y , $A\hat{T}T$, is given by:

$$A\hat{T}T = \bar{y}_1 - \frac{1}{N} \sum_{i=1}^N \left[\frac{d_i \hat{y}_i^0}{\hat{P}} + \frac{(1 - d_i) \hat{\omega}_i (y_i - \hat{y}_i^0)}{\hat{P}} \right] \quad (5)$$

$$\hat{\omega}_i = \frac{\hat{p}(\mathbf{x}_i)}{1 - \hat{p}(\mathbf{x}_i)}$$

where $\bar{y}_1$ is the sample average of the outcome among RBD participants, $\hat{P}$ is the proportion of the sample participating in the RBD program, $\hat{p}(\mathbf{x}_i)$ is the estimated propensity score, and $\hat{y}_i^0$ is the value of the outcome in the absence of RBD participation as predicted by OLS or the fitted probability as predicted by a logit regression, depending

on whether the outcome is continuous or binary. As shown in the appendix, the ATT estimator in (5) is “doubly robust”. That is, the ATT estimator is unbiased when either the model for the outcome or that of the propensity score is correctly specified (Wooldridge 2007).

Note that nearly all of our outcome variables are nonnegative and continuous. In order to ensure that modeling each outcome by OLS is appropriate, we apply the inverse hyperbolic sine (IHS) transformation (Burbidge, Magee and Robb 1988) to all nonnegative continuous outcomes prior to variable selection and ATT estimation:

$$\ln(y_i + \sqrt{y_i^2 + 1}) \quad (6)$$

By virtue of applying the IHS transformation, the coefficients from the linear regression model of a given outcome can be interpreted as if the dependent variable were in log form.

4.3 Inference

In the inference stage, we create t -statistics using the square root of the variance estimator derived by Farrell (2015):

$$\frac{1}{N} \left[\sum_{i=1}^N \frac{d_i(y_i - \hat{y}_i^0 - A\hat{T}T)^2}{\hat{P}^2} + \frac{(1 - d_i)\hat{\omega}_i(y_i - \hat{y}_i^0)^2}{\hat{P}(1 - \hat{P})} \right] \quad (7)$$

The ATT estimator is asymptotically normal and the variance estimator is robust to heteroscedasticity of unknown form. Farrell (2015) uses critical values from the standard normal distribution in hypothesis testing. In order to make inference appropriately conservative given our small sample, we use critical values from a t distribution with

degrees of freedom equal to the total number of households in the data set minus the total number of parameters in the propensity score and outcome models.

We adjust inference for multiple hypothesis testing by controlling for the proportion of rejected null hypotheses that are true, or the “False Discovery Rate” (FDR), by reporting the “ q -value” of each hypothesis test as calculated using the method of Benjamini and Hochberg (1995). To understand what is measured by a given q -value, suppose that a set of hypothesis test is organized into “families”, where each family consists of tests providing evidence for a particular question, e.g. whether the RBD program affected plantain production. Suppose a test of a particular hypothesis in the family generates a p -value of 0.06 and a q -value of 0.10. The p -value is interpreted in the usual way. The q -value indicates that if we were to reject all hypotheses in the family with p -values no greater than 0.06, no more than 10% of all rejected hypotheses would be true in expectation. In other words, the FDR associated with using a cutoff rule of $p < 0.06$ for hypothesis testing is no greater than 0.10. Each q -value is calculated as $p \times M / r$, where M is the number of hypothesis tests in the relevant family of hypotheses and r is the rank of the p -value within its family when ordered from smallest to largest.

We divide outcomes into families of hypotheses that provide evidence for whether the RBD program affected household consumption expenditure (with outcome *Expenditure per adult male equivalent*), plantain revenue (*Revenue* and *Revenue per manzana*), plantain production (*Plantain production*, *Plantain yield*, *Commercial quality plantain production*, and *Commercial quality plantain yield*), irrigated area (*Irrigated area producing plantain* and *Irrigated plantain area*), land use (*Area producing plantain*, *Area planted with*

plantain, and *Area in maize*), exposure to production shocks (binary indicators for *Plantain affected by drought* and *Plantain affected by other shocks*), value of inputs applied to plantain (*Value of fertilizer applied to plantain*, *Value of chemicals applied to plantain*, *Fertilizer value per manzana*, *Chemical value per manzana*) or credit received (*Total credit*).

Since the mechanisms driving RBD program impacts were quite different across years (with much stronger irrigation effects in year two), we attempt to answer the above questions separately for year one and year two of the program. Note that for irrigated area and area in plantain, we distinguish between area planted in the moment of the interview and area producing plantain since irrigation was installed on beneficiary plots near the end of the first year of the program. For production, we distinguish between first grade or “commercial quality” plantain appropriate for sale to large food retailers and second and third grade plantain.

4.4 Robustness checks

Our primary robustness check comes from Athey and Imbens (2015) and Athey et al. (2017). For each covariate (not including higher-order terms), we split the sample at the median of the covariate and estimate all ATTs on the resulting subsamples. We repeat the entire variable selection and estimation process, without the split covariate, on each subsample: generation of second order terms, model selection by the LASSO and cross validation, and doubly robust estimation. We then compute a new ATT estimate by taking a weighted average of the two subsample ATTs, where the weights are given by each subsample’s share of the overall sample.

We repeat this process for all covariates until we have 23 new estimates of each ATT, one for each split covariate. There is no accepted heuristic (e.g. $p < 0.05$ in the case of hypothesis testing) that can be applied to the results of this procedure in deciding whether our results appear to be robust. However, the standard deviation of the resulting ATT estimates for a given outcome is informative with respect to the sensitivity of our results to model specification. A small standard deviation relative to the standard error of the full sample ATT would suggest that sampling variability swamps any uncertainty attributable to model specification, and that our results are quite robust. On the other hand, if the ratio of the standard error of the full ATT to the standard deviation of the split-sample ATTs is small, then our results are probably not robust, and the estimated standard error might not provide an adequate measure of uncertainty. In addition, by comparing our full-sample ATT estimates to the average and standard deviation of the split-sample ATT estimates, we can see whether our full-sample ATT estimates are outliers relative to other estimated ATTs we could have obtained with our covariates.

5. Results

We begin by checking whether inverse probability weighting improves balance for covariates chosen by our model selection procedure.

5.1 Covariate balance with and without inverse probability weighting

The logit LASSO and cross validation procedure for the propensity score model selected the five covariates listed below the linearized propensity score in the first column of table 2. The linearized propensity score is given by the log odds ratio, $\ln[\hat{p}(\mathbf{x}_i) / (1 - \hat{p}(\mathbf{x}_i))]$.

The next three columns show the normalized difference (equal to $|\bar{x}_1 - \bar{x}_0| / \sqrt{0.5(s_1^2 + s_0^2)}$, where $\bar{x}_d$ and s_d^2 are the average and standard deviation for the subsample with treatment status d , respectively), the p -value from a test of equality of means, and the ratio of the standard deviation among participant households to the standard deviation among non-participants when assigning equal weight to each observation in the data. The last three columns report the same results for the inverse probability weighted sample.

[TABLE 2 HERE]

The linearized propensity score provides a scalar summary of the joint distribution of the covariates. Prior to inverse probability weighting, the difference in means for the linearized score is statistically significant (p -value < 0.001) while the standard deviation of the linearized score is very similar in both subsamples. After inverse probability weighting, the normalized difference for the linearized score is nearly zero and is no longer statistically significant. The ratio of standard deviations for the linearized score in the two subsamples is slightly improved with inverse probability weighting as well. For the individual covariates, the normalized differences are sharply reduced after weighting, while the ratio of standard deviations is much closer to unity for two of the covariates and barely changed for the remaining three covariates after weighting. The inverse probability weights are effective in improving balance for covariates that are associated with RBD participation.

5.2 Estimated Average Treatment Effects on the Treated, year one

The estimated ATTs for the first year of RBD participation are presented in table 3. *Plantain affected by drought* and *Plantain affected by other shocks* are binary indicators, with the former equal to 1 if the household indicated that any plantain plots were affected by drought and the latter equal to 1 if the household stated that any plantain plots were affected by pests, flooding, plant viruses, and other non-drought shocks. The ATTs for the shock indicators are impacts of the program on the probability of reporting the relevant type of shock. We distinguish between drought and other shocks since irrigation might be expected to have stronger effects on the incidence of drought. The remaining ATTs describe average impacts on the natural log of each outcome. Alongside the ATTs we report percent impacts, which give the ATT as a percent of the average value of the outcome that would have been observed among RBD participants had they not enrolled in the program. For outcomes in natural log form, we estimate the percent changes as $100 \times (\exp(A\hat{T}T) - 1)$. For binary outcomes, the percent changes are given by $100 \times A\hat{T}T / (\bar{y}_1 - A\hat{T}T)$.

[TABLE 3 HERE]

There is no evidence for an impact on household expenditure in year one. Although the point estimates of the ATTs for some outcomes related to revenue and production are large, they are imprecisely estimated, with 95% confidence intervals that are consistent with very small effects. The ATTs obtained through sample splitting suggest that impact estimates for outcomes related to revenue and production are somewhat sensitive to model

specification. In some cases, the standard deviations of the split-sample ATTs are quite large while some of the averages are much smaller than the initial point estimates. Overall, we conclude that there is no strong evidence for or against RBD program effects in year one on revenue or production.

The ATT for *Irrigated plantain area* is large and precisely estimated, while the ATT for *Irrigated area producing plantain* is small and not statistically significant. Recall that irrigation was not installed until late in the first year of participation. The size of the effect for *Irrigated area producing plantain* suggests that irrigation was installed on land with new plantain, an observation that is consistent with the program's promotion of improved planting material. However, the small and insignificant impacts on *Area producing plantain* and *Area planted with plantain* suggest that while the program appears to have succeeded in expanding irrigation, much of the land receiving irrigation and improved planting material would have produced plantain regardless of RBD participation status, at least in year one.

The estimated ATTs for the two shock indicators suggest that the program succeeded in reducing exposure to drought in year one. While the impact on reported drought shocks is large in relative terms, the 95% confidence interval on the effect is consistent with an impact as small as 2%. As shown in the split-sample results, the drought effect is robust to changes in model specification. It is worth noting that the 2009-2010 agricultural year was a time of widespread drought in Central America (El Mundo 2010). Inputs, planting material, and extension accessed by RBD participants may have helped make plantain more robust to stress from a lack of moisture.

The signs and magnitudes of the impacts on input value suggest that program subsidies increased use of chemicals while having mixed effects on fertilizer. While the effect on *Fertilizer value per manzana* is statistically significant, its *q*-value suggests that the effect may have arisen by chance. The ATT for *Total credit* is large in absolute value, and suggests that on average beneficiaries received \$1,829 less in credit than they would have in the absence of the RBD program. The estimated average reduction in credit is equal to about 29% of the per-beneficiary subsidy paid by MCC. However, the 95% confidence interval for the ATT on *Total credit* is consistent with small positive effects as well as large negative effects, and the ATT is an extreme negative outlier relative to the alternative estimates obtained through sample splitting. We conclude that in year one, there is strong evidence for a positive effect of the RBD program on irrigated plantain area, evidence for a reduction in the effects of drought, and no clear evidence for other effects.

5.3 Estimated Average Treatment Effects on the Treated, year two

The estimated ATTs for year two of participation in the RBD program are shown in table 4. As in year one, there is no discernible effect on household expenditure. Impacts on revenue and production indicators are large and statistically significant. In interpreting the large estimated ATTs, it should be kept in mind that the 95% confidence intervals for some outcomes are consistent with relatively modest impacts. The intervals around *Revenue*, *Revenue per manzana*, *Plantain yield*, and *Commercial quality plantain yield* are consistent with effects as small as 21%, 7%, -3%, and 4%, respectively. Despite these caveats, the evidence clearly suggests that the RBD program had a strong effect on revenue and

production, both overall and per manzana, once irrigation was in place. The full-sample ATTs for *Revenue* and *Plantain production* are large relative to the ATTs obtained through sample splitting. That being said, the standard deviations for the split-sample ATTs are quite small. Even the extreme left tail of the distribution of split-sample ATTs contains estimated impacts that are quite large for all revenue and production indicators.

[TABLE 4 HERE]

Dividing the observed averages of *Area producing plantain* and *Irrigated area producing plantain* among RBD participants in year two by one plus their respective percent impacts indicates that in the absence of the program, RBD participants would have been producing plantain on an average of 1.63 manzanas, 66% of which would have been irrigated. After having participated in the RBD program, participants produce plantain on an average of 2.04 manzanas, 93% of which are irrigated.

While *Area producing plantain* shows an impact of about 24%, *Area planted with plantain* only shows an impact of 11%, and the ATT is not statistically significant. These results suggests that were we to have conducted a third survey round, the impacts on *Area producing plantain* might be smaller than what is shown in table 4. But it should be noted that the estimated ATT for *Area planted with plantain* is an outlier relative to the split-sample ATTs. The estimates that are two standard deviations to the left of the mean of the split-sample ATTs are no smaller than 0.18, or an impact of about 20%. While it appears that the expansion of irrigated area was not matched one-for-one by growth in planted area,

the RBD program does appear to have increased area producing plantain while achieving nearly complete irrigation coverage.

The estimated ATTs offer no evidence that the RBD program affected self-reported shocks in year two, although our robustness checks suggest that these might be underestimates of the program impact on shocks. In contrast to year one, the principal climatic shock in year two came in the form of excessive rain from Hurricane Matthew (El Nuevo Diario 2010), and it is not clear how extension, inputs, and irrigation would help offset flooding or storm damage. While the point estimate for *Area in maize* suggests that RBD beneficiaries may have reduced maize area in order to make room for plantain, the effect is imprecisely estimated. Impacts on input use are not reported for year two, as they were only collected in year one. However, administrative data indicate that the value of inputs applied through RBD participation did not change much from year to year, amounting to \$1,515 and \$1,624 per beneficiary, respectively, in years one and two for the cohort studied in this paper. Therefore year one impacts on input value may be representative of impacts in year two.

As in the case of year one, the impact on credit received is large and negative, although it is not statistically significant. Unlike in the case of year one, the split-sample results indicate that the estimated ATT for *Total credit* in year two is robust to changes in model specification. The estimated percent impact of -60% suggests that credit received by beneficiaries fell by \$953 as a result of RBD participation, or about 15% of the per beneficiary subsidy paid by MCC. We will return to the results for *Total credit* when

examining possible explanations for the absence of impacts on *Expenditure per adult male equivalent* in the conclusion to this article.

5.4 Heterogeneous treatment effects

As a byproduct of our robustness check, we obtain estimates of the ATT for each subgroup defined by splitting the covariates at their respective medians. Overall, there is clear evidence that households below the median for land area, indices of agricultural and non-agricultural wealth, and credit received at baseline experienced large impacts on irrigated area, planted area, and input value relative to other farmers in the sample. Many of the subgroup comparisons for irrigation and land use variables are statistically significant, and a large number of the estimated subgroup ATTs have low q -values when adjusting for the number of subgroup ATTs estimated within each outcome variable. Note that splitting the sample at the median of pre-program credit we divide the sample into farmers with and without loans prior to RBD participation. While the full set of estimated subgroup effects is presented in the supplementary appendix, the subgroup impacts on plantain land use and irrigation are summarized in figure 1.

[FIGURE 1 HERE]

There is also a pattern of larger impacts on outcomes related to production and revenue for farmers that are poorer and have no credit prior to RBD participation. However, unlike the set of outcomes related to irrigation and planted area, many of the estimated subgroup ATTs appear to have arisen by chance after adjusting for multiple hypothesis

testing. Figure 2 shows the ATTs for outcomes related to sales and production in year two, when all program components were in place.

[FIGURE 2 HERE]

The subgroup treatment effects suggest that the RBD program was particularly effective in increasing investment among farmers who are relatively poor and had no loans prior to program enrollment. To the extent that this group of farmers would not have had the resources necessary for expanding irrigation and planted area in the absence of the RBD program, the matching funds component of the RBD program appears to have been successful. The corresponding impacts on sales, production, and yield show less heterogeneity by wealth and credit access. While the data offer no clear answer for the difference in heterogeneity of effects across outcome categories, it may be the case that the RBD program was effective in loosening liquidity constraints on relatively poor farmers but less able to address other constraints that might hinder performance.

5.5 Additional robustness checks

In the supplementary appendix, we present the results of three additional robustness checks. First, we repeat our analysis using only the observations on the common support of the propensity score. As shown in tables 3, 4, and 5 of the appendix, restricting the sample to the common support did not improve covariate balance, forced us to drop nine treated observations (thereby changing the interpretation of the effect being estimated), and generated results that were nearly identical to results for the full sample. Therefore we opted to include results for the full sample in the main text.

Second, for outcomes that are observed pre-RBD as a result of survey recall questions (*Irrigated plantain area*, *Area planted with plantain*, *Total credit*, and *Area in maize*), we estimate each ATT by difference-in-differences (DID). DID results are presented in table 5. The signs and statistical significance of the estimated DID effects are all the same as our estimates obtained under the CIA. The effect on *Irrigated plantain area* is substantially larger when estimated by DID relative to our main results, reinforcing our conclusion that the program had a large effect on irrigation. The impact on *Total credit* in year one is estimated much less precisely by DID relative to what was obtained in our results, suggesting that the RBD program may not have crowded out credit in year one. All other DID estimates are quite similar to our main results. We conclude that our results are robust to changes in our identifying assumptions.

[TABLE 5 HERE]

Lastly, for the outcomes observed pre-program, we estimate “placebo” ATTs. That is, we estimate ATTs for the year prior to the start of the RBD program for the cohort studied in this paper. When estimating placebo ATTs (as well as DID effects), we cannot use any pre-intervention outcomes as control variables. Therefore we expect the placebo results to understate the degree to which our empirical strategy is able to address selection bias. Placebo ATT results are shown in table 6. All placebo effects are small or imprecisely estimated. When combined with the fact that our ability to control bias when estimating placebo ATTs is reduced, the results in table 6 support our identification strategy.

[TABLE 6 HERE]

6. Conclusion

Using a combination of ML methods and doubly robust estimation, we assessed the impacts of the RBD program for smallholder plantain producers. In year one, program beneficiaries received subsidies for extension, seed, and inputs, while having irrigation installed late in the agricultural year. In year two, beneficiaries once again received extension and inputs, while also enjoying the full effects of expanded irrigation. Evidence for program impacts in year one, when beneficiaries could not have experienced anything but trivial impacts of irrigation, are mixed. In year one we see large impacts on irrigated area and a reduction in self-reported effects of drought, but no clear evidence of impacts on expenditure, revenue, production, yield, planted area, input value, self-reported plantain shocks, or credit. In year two, we see strong evidence for large effects on revenue, production, yield, and irrigated area, with mixed evidence for effects on area in plantain.

Our subgroup analysis indicates that the program was particularly effective in expanding planted area and irrigation among farmers who were poor and lacked access to credit prior to enrollment in the RBD program. These same farmers appear to also have enjoyed relatively large impacts on revenue, production, and yield, but the evidence is less clear cut than it is for irrigation and planted area.

The absence of any detectable impacts on household consumption expenditure is surprising given the large estimated increases in revenue. The estimated ATTs point to two possible explanations. First, the estimated ATT for *Area in maize* in year two suggests that beneficiaries reduced maize area by 19% as a result of RBD participation. Maize prices in Nicaragua were quite high in 2009 and 2010 (FAO 2017). The long-term investment made

by beneficiaries in the perennial crop of plantain may have come at the short-term opportunity cost of foregone maize revenue. That being said, the 95% confidence interval for *Area in maize* is consistent with small positive effects (as large as 7%) as well as large negative effects, and the results obtained by splitting the sample suggest that the full-sample ATT may be implausibly large in absolute value. Second, the negative ATT for *Total credit* in year two suggests that RBD subsidies may have partly substituted for other credit sources. While results for *Total credit* in year two are consistent with small positive effects as well as large negative effects, the sign and magnitude of the point estimate are robust to changes in identifying assumptions (i.e. moving from CIA to the assumptions underlying DID) and model specification (as shown by the split-sample results). There appears to be mild evidence for RBD impacts being diminished by the crowding out of credit.

A third potential explanation for the lack of household expenditure growth is that while revenue impacts were large, the program was also quite costly to beneficiaries. The estimated ATT on the natural log of revenue suggests that plantain revenue grew by \$832 in year one (although the impact is not statistically significant) and \$4,598 in year two as a result of the program. While comprehensive data on production costs are not available for the full sample of farmers in our data, we do have detailed administrative cost data for all participants. A simple pre-post estimate of the impact of the RBD program on total cost shows that variable costs (labor, land rent, fertilizer, and chemicals) increased from \$809 at baseline to \$4,213 per producer in year one. Note that this does not include the portion of costs paid for by the program. Participants also spent \$1,703 each on irrigation equipment and \$850 on land preparation in year one. In year two, participants spent an

average of \$4,040 on variable inputs. In the absence of a counterfactual estimated with a control group, it is hard to know whether these back-of-the-envelope estimates for cost impacts are credible. But the observed growth in production costs suggests that large impacts on sales may have corresponded to more modest effects on profits.

Other studies of smallholder matching funds programs have found positive effects on the adoption of technologies or production practices (Ton et al. 2015; Maffioli et al. 2013) as well as production, yield, and revenue (Ton et al. 2015; Carter, Tjernström and Toledo 2016; Mullally and Maffioli 2016). Our results support the positive findings of the existing literature for the indicators just described. However, as in the case of previous studies, we fail to find welfare impacts of the RBD program. The high cost of land preparation and infrastructure for irrigation suggest that RBD beneficiaries may be foregoing consumption gains in the short term in order to invest in their farms, in which case expenditure gains would not be realized until years later. But without additional rounds of survey data, we can only speculate as to how RBD impacts might evolve going forward.

The institutional burden of groundwater irrigation is low in that cooperation among farmers is not required for its development, although resources can be over-exploited in the absence of collective management. In contrast, surface water projects require the creation of water-user groups and the maintenance of shared infrastructure. In this sense, groundwater projects like the RBD plantain program appear ideal for developing countries with weak institutions. But the fixed costs of shared surface water infrastructure can be spread over a large number of producers, potentially driving down average cost of

production as the number of smallholder beneficiaries expands. The suitability of different irrigation investments in a given area will depend not only the quality of local institutions, but also on the existing scale of production among farmers. Carefully examining how the costs and benefits of groundwater irrigation projects vary with farm size is an area for future research.

References

- Adeoti, A., B. Barry, R. Namara, and A. Kamara. 2009. "The Impact of Treadle Pump Irrigation Technology Adoption on Poverty in Ghana." *The Journal of Agricultural Education and Extension* 15 (4): 357-369.
- Anderson, M. 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.
- Angrist, J., and J. Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton, New Jersey: Princeton Press.
- Athey, S., and G. Imbens. 2015. "Machine Learning Methods in Economics and Econometrics: A Measure of Robustness to Misspecification." *American Economic Review: Papers & Proceedings* 105 (5): 476-480.
- Athey, S., G. Imbens, T. Pham, and S. Wager. 2017. "Estimating Average Treatment Effects: Supplementary Analyses and Remaining Challenges." *American Economic Review: Papers & Proceedings* 107 (5): 278-281.
- Barrett, C., M. Bachke, M. Bellemare, H. Michelson, S. Narayanan, and T. Walker. 2012. "Smallholder Participation in Agricultural Value Chains: Comparative Evidence from Three Countries." *World Development* 40 (4): 715-730.

- Belloni, A., V. Chernozhukov, and C. Hansen. 2014. "High-Dimensional Methods and Inference on Structural and Treatment Effects." *Journal of Economic Perspectives* 28 (2): 29-50.
- Belloni, A., V. Chernozhukov, and C. Hansen. 2014. "Inference on Treatment Effects after Selection among High-Dimensional Controls." *Review of Economic Studies* 81 (2): 605-650.
- Benjamini, Y., and Y. Hochberg. 1995. "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing." *Journal of the Royal Statistical Society* 57 (1): 289-300.
- Burbidge, J., L. Magee, and L. Robb. 1988. "Alternative Transformation to Handle Extreme Values of the Dependent Variable." *Journal of the American Statistical Association* 83 (401): 123-127.
- Carter, M., P. Toledo, and E. Tjernström. 2016. *Heterogeneous Impact Dynamics of a Rural Business Development Program in Nicaragua*. Working Paper, Cambridge, MA: National Bureau of Economic Research. Accessed July 1, 2017. <http://www.nber.org/papers/w22628.pdf>.
- Cattaneo, M. 2010. "Efficient Semiparametric Estimation of Multi-Valued Treatment Effects under Ignorability." *Journal of Econometrics* 155 (2): 138-154.
- Crespi, G., and A. Maffioli. 2014. "Design and Evaluation of Fiscal Incentives for Business Innovation in Latin America: Lessons Learned after 20 Years of Experimentation."

- In *Science, Technology and Innovation Policies for Development*, edited by G. Crespi and G. Dutrenit. New York: Springer International.
- Danilov, D., and J. Magnus. 2004. "On Ignoring the Harm that Pretesting Can Cause." *Journal of Econometrics* 122 (1): 27-46.
- Deaton, A., and S. Zaidi. 2002. *Guidelines for Constructing Consumption Aggregates for Welfare Analysis*. Living Standards Measurement Study Working Paper, Washington, DC: World Bank.
- Del Carpio, X., N. Loayza, and G. Datar. 2011. "Is Irrigation Rehabilitation Good for Poor Farmers? An Impact Evaluation of a Non-Experimental Irrigation Project in Peru." *Journal of Agricultural Economics* 62 (2): 449-473.
- Dillon, A. 2011a. "Do Differences in the Scale of Irrigation Projects Generate Different Impacts on Poverty and Production?" *Journal of Agricultural Economics* 62 (2): 474-492.
- Dillon, A. 2011b. "The Effect of Irrigation on Poverty Reduction, Asset Accumulation, and Informal Insurance: Evidence from Northern Mali." *World Development* 39 (12): 2165-2175.
- Doss, C. 2014. "Data Needs for Gender Analysis in Agriculture." In *Gender in Agriculture*, edited by A. Quisumbing, R. Meinzen-Dick, T. Raney, A. Croppenstedt, J. Behrman and A. Peterman, 55-68. Dordrecht: Springer Science.

- Edwards, E., and S. Smith. 2017. "The Role of Irrigation in the Development of Agriculture in the United States." Unpublished manuscript.
- El Mundo. 2010. "The Drought in Northern Nicaragua Requires Food Distribution (La Sequía del Norte de Nicaragua Obliga la Distribución de Alimentos)." *El Mundo*, March 23. Accessed November 1, 2017. <http://www.elmundo.es/america/2010/03/23/noticias/1269384568.html>.
- El Nuevo Diario. 2010. "Devastating Damages (Daños Devastadores)." *El Nuevo Diario*, September 28. Accessed November 5, 2017. <https://redesdesolidaridad.wordpress.com/2010/09/29/alarma-por-lluvias-en-nicaragua-danos-evacuaciones-y-damnificados/>.
- Fan, S., A. Gulati, and S. Thorat. 2008. "Investment, Subsidies, and Pro-Poor Growth in Rural India." *Agricultural Economics* 39 (2): 163-170.
- FAO. 2017. *FAOSTAT*. Accessed 06 30, 2017. <http://www.fao.org/faostat/>.
- Farrell, M. 2015. "Robust Inference on Average Treatment Effects with Possibly More Covariates than Observations." *Journal of Econometrics* 189 (1): 1-23.
- Foster, A., and M. Rosenzweig. 2010. "Microeconomics of Technology Adoption." *Annual Review of Economics* 2 (1): 395-424.
- Hagos, F., G. Jayasinghe, S. Awulachew, M. Loulseged, and A. Yilma. 2012. "Agricultural Water Management and Poverty in Ethiopia." *Agricultural Economics* 43 (1): 99-111.

- Hastie, T., R. Tibshirani, and M. Wainwright. 2015. *Statistical Learning with Sparsity: The Lasso and Generalizations*. Boca Raton, FL: CRC Press.
- Huang, Q., S. Rozelle, B. Lohmar, J. Huang, and J. Wang. 2006. "Irrigation, Agricultural Performance and Poverty Reduction in China." *Food Policy* 31 (1): 30-52.
- Imbens, G. 2004. "Nonparametric Estimation of Average Treatment Effects under Exogeneity: A Review." *The Review of Economics and Statistics* 86 (1): 4-29.
- INCAP. 2013. *Complimentary Study for the Analisis of Data from the Guatemala National Living Standards Survey (Estudio Complementario al Analisis Secundario de los Datos de la Encuesta Nacional de Condiciones de Vida de Guatemala)*. Guatemala City: INCAP.
- Kuwornu, J., and E. Owusu. 2012. "Irrigation Access and Per Capita Consumption Expenditure in Farm Households: Evidence from Ghana." *Journal of Development and Agricultural Economics* 4 (3): 78-92.
- Nankhuni, F., and G. Paniagua. 2013. *Meta-Evaluation of Private Sector Interventions in Agribusiness*. Washington, DC: World Bank.
- Palanisami, K. 1997. "Economics of Irrigation Technology Transfer and Adoption." *Irrigation Technology Transfer in Support of Food Security*. Rome: FAO.
- Ton, G., K. de Grip, L. Klerkx, M. Rau, M. Douma, E. Friis-Hansen, B. Triomphe, A. Waters-Bayer, and M. Wongtschowski. 2013. *Effectiveness of Innovation Grants to*

- Smallholder Agricultural Producers: An Explorative Systematic Review*. London: University of London.
- Van den Berg, M., and R. Ruben. 2006. "Small-Scale Irrigation and Income Distribution in Ethiopia." *Journal of Development Studies* 42 (5): 868-880.
- Wooldridge, J. 2007. "Inverse Probability Weighted Estimation for General Missing Data Problems." *Journal of Econometrics* 141 (2): 128-1301.
- Zhu, M., and A. Ghodsi. 2006. "Automatic Dimensionality Selection from the Scree Plot via the Use of Profile Likelihood." *Computational Statistics & Data Analysis* 51 (2): 918-930.

Figure 1. Heterogeneous treatment effects for plantain land use, irrigation, and input value

Notes: The figure shows the kernel density plots of the subsample ATTs in years one and two for *Irrigated plantain area*, *Irrigated area producing plantain*, *Area planted with plantain*, *Area producing plantain*, *Value of fertilizer applied to plantain*, *Value of chemicals applied to plantain*, *Fertilizer value per manzana*, *Chemical value per manzana*, and *Total credit received*. “Above median covariate value” refers to households above the pre-program median for *Agricultural wealth factor 1*, *Agricultural wealth factor 2*, *Agricultural wealth factor 3*, *Land area*, *Number of agricultural parcels*, *Total credit*, *Non-agricultural wealth factor 1*, or *Non-agricultural wealth factor 2*.

Figure 2. Heterogeneous treatment effects for plantain revenue and production outcomes

Notes: The figure shows the kernel density plots of the subsample ATTs in year two for *Revenue*, *Revenue per manzana*, *Plantain production*, *Plantain yield*, *Commercial quality plantain production*, and *Commercial quality plantain yield*. “Above median covariate value” refers to households above the pre-program median for *Agricultural wealth factor 1*, *Agricultural wealth factor 2*, *Agricultural wealth factor 3*, *Land area*, *Number of agricultural parcels*, *Total credit*, *Non-agricultural wealth factor 1*, or *Non-agricultural wealth factor 2*.

Table 1. Summary Statistics, Pre-RBD Program

	<u>RBD participants</u>		<u>Non-participants</u>		<i>p</i> -value for difference in means
	Average	Standard deviation	Average	Standard deviation	
Female household head (0/1)	0.200	0.402	0.157	0.367	0.525
Household size (adult equivalents)	1.690	0.522	1.754	0.499	0.474
Years of education of most educated member	9.653	4.153	9.275	3.914	0.594
Average age of adults in household (years)	35.971	10.635	36.909	10.915	0.616
Owens a non-farm business (0/1)	0.179	0.385	0.353	0.483	0.019
Belongs to a producer cooperative (0/1)	0.253	0.437	0.216	0.415	0.621
Farm located in municipality 1 (0/1)	0.105	0.309	0.078	0.272	0.603
Land area (manzanas)	1.949	0.920	1.552	0.946	0.015
Number of agricultural parcels	1.421	0.894	1.235	0.737	0.206
Agricultural wealth factor 1	-0.069	1.727	-0.020	2.061	0.880
Non-agricultural wealth factor 1	0.007	2.658	-0.027	2.047	0.937
Total credit (dollars)	965.969	2,755.827	1,243.549	3,226.809	0.586
Unearned income (dollars)	143.366	748.003	100.686	432.849	0.708
Years in last ten with plantain	4.021	3.028	4.137	3.231	0.829
Simpson index of crop diversity	0.359	0.404	0.344	0.433	0.833
Area in maize (manzanas)	0.422	0.660	0.257	0.565	0.134
Area planted with plantain (manzanas)	0.455	0.578	0.539	0.576	0.408
Irrigated area in plantain (manzanas)	0.278	0.951	0.366	1.025	0.605
Observations	95		51		

Notes: Adult equivalents calculated using weights for Central America from INCAP (2013). *Agricultural wealth factor 1* and *Non-agricultural wealth factor 1* are the first factors from principal components decompositions of agricultural and non-agricultural wealth variables, respectively. Agricultural wealth variables include agricultural installations, tools and mobile equipment, and livestock. Non-agricultural wealth variables include consumer durables and access to running water, trash collection, and sanitation. The data set contains households located in two distinct municipalities, and “municipality 1” is a dummy variable equal to 1 for all households in one of the two municipalities.

Table 2. Covariate balance with and without inverse probability weighting

	Unweighted sample			Inverse probability weighted sample		
	Normalized difference	<i>p</i> -value	Ratio of standard deviations	Normalized difference	<i>p</i> -value	Ratio of standard deviations
Linearized propensity score	0.90	0.000	0.91	0.01	0.960	1.05
Years of education $\times$ Land area	0.44	0.015	1.24	0.14	0.539	1.27
Land area	0.43	0.015	0.97	0.02	0.935	1.05
Number of agricultural parcels $\times$ Owns a non-farm business	0.41	0.011	0.60	0.13	0.374	1.08
Total credit $\times$ Agricultural wealth factor 2	0.34	0.027	0.41	0.13	0.600	0.78
Area planted with plantain $\times$ Owns a non-farm business	0.46	0.006	0.72	0.00	0.988	1.30
Observations		146			146	

Notes: Propensity score estimated via a logit regression, with covariates selected using the logit LASSO. The LASSO tuning parameter was selected by five-fold cross validation as applied to the post-LASSO (i.e. unpenalized) model. The normalized differences show the difference in means divided by the pooled sample standard deviation as described in the text.

Table 3. Estimated Average Treatment Effects on the Treated, RBD Program Year 1

	Full sample results					Split sample results	
	ATT	Percent impact	Standard error	p-value	q-value	Mean ATT	Standard deviation
Expenditure per adult male equivalent	0.08	8.16%	0.07	0.288	0.288	0.16	0.03
Revenue	0.34	40.91%	0.48	0.475	0.950	0.26	0.28
Revenue per manzana	0.15	15.83%	0.49	0.767	0.767	0.10	0.26
Plantain production	0.32	37.67%	0.26	0.218	0.870	0.22	0.16
Plantain yield	0.13	13.60%	0.27	0.640	0.640	0.08	0.14
Commercial quality plantain production	0.28	32.04%	0.25	0.264	0.528	0.14	0.15
Commercial quality plantain yield	0.20	22.50%	0.27	0.447	0.447	0.01	0.13
Irrigated plantain area	0.42	51.78%	0.09	0.000	0.000	0.37	0.06
Irrigated area producing plantain	0.01	0.64%	0.09	0.940	0.940	0.14	0.05
Area producing plantain	0.08	8.69%	0.06	0.149	0.448	0.10	0.05
Area planted with plantain	0.07	7.09%	0.10	0.504	0.504	0.27	0.05
Area in maize	0.07	7.14%	0.08	0.365	0.547	0.17	0.02
Plantain affected by drought	-0.19	-41.67%	0.08	0.027	0.055	-0.12	0.02
Plantain affected by other shocks	-0.10	-15.93%	0.09	0.272	0.272	-0.04	0.03
Value of fertilizer applied to plantain	0.06	6.28%	0.15	0.681	0.681	-0.02	0.06
Value of chemicals applied to plantain	0.35	41.68%	0.32	0.273	0.547	0.33	0.14
Fertilizer value per manzana	-0.40	-32.70%	0.19	0.037	0.150	-0.37	0.03
Chemical value per manzana	0.25	28.93%	0.31	0.413	0.550	0.04	0.10
Total credit	-1.02	-64.00%	0.59	0.083	0.083	0.03	0.19
Observations	146						

Notes: *Expenditure per adult male equivalent* and *Value of fertilizer applied to plantain* are in natural log form. *Plantain affected by drought* and *Plantain affected by other shocks* are binary indicators. The IHS transformation was applied to all other dependent variables. ATT and standard errors estimated as in Farrell (2015). The percent impacts give the average percent increase in the outcome among RBD beneficiaries relative to the counterfactual “no RBD program” average outcome. All hypothesis tests conducted using a *t*-distribution with degrees of freedom given by the sample size minus the total number of parameters in the propensity score and regression models for each ATT. “Median split” results were calculated by splitting each original covariate at its median, re-estimating the ATT on each subsample, and generating a new estimate of the ATT using a weighted average of the subsample estimates. The median split results are the mean and standard deviation of the resulting ATT estimates. For each ATT, the *q*-value gives the proportion of rejected null hypotheses that would in fact be true if we rejected all null hypotheses with *p*-values no greater than what is reported for the given ATT. The *q*-values were calculated using the method of Benjamini and Hochberg (1995).

Table 4. Estimated Average Treatment Effects on the Treated, RBD Program Year 2

	ATT	Full sample results				Split sample results	
		Percent impact	Standard error	<i>p</i> -value	<i>q</i> -value	Mean ATT	Standard deviation
Expenditure per adult male equivalent	-0.04	-4.01%	0.07	0.574	0.574	0.01	0.03
Revenue	1.25	249.12%	0.53	0.021	0.042	0.98	0.06
Revenue per manzana	0.74	110.47%	0.34	0.030	0.030	0.72	0.05
Plantain production	0.83	129.54%	0.26	0.001	0.006	0.61	0.05
Plantain yield	0.38	45.65%	0.20	0.068	0.068	0.35	0.04
Commercial quality plantain production	0.81	125.02%	0.27	0.003	0.005	0.60	0.07
Commercial quality plantain yield	0.48	61.47%	0.22	0.033	0.045	0.40	0.06
Irrigated plantain area	0.54	72.10%	0.11	0.000	0.000	0.60	0.06
Irrigated area producing plantain	0.55	73.21%	0.10	0.000	0.000	0.51	0.05
Area producing plantain	0.21	23.52%	0.09	0.019	0.057	0.18	0.03
Area planted with plantain	0.11	11.30%	0.09	0.261	0.261	0.28	0.05
Area in maize	-0.21	-18.68%	0.14	0.145	0.217	0.03	0.02
Plantain affected by drought	0.00	0.73%	0.08	0.979	0.979	-0.13	0.01
Plantain affected by other shocks	0.11	16.33%	0.08	0.172	0.343	-0.04	0.02
Total credit	-0.93	-60.53%	0.57	0.103	0.103	-0.89	0.11
Observations	146						

Notes: *Expenditure per adult male equivalent* is in natural log form. *Plantain affected by drought* and *Plantain affected by other shocks* are binary indicators. The IHS transformation was applied to all other dependent variables. ATT and standard errors estimated as in Farrell (2015). The percent impacts give the average percent increase in the outcome among RBD beneficiaries relative to the counterfactual “no RBD program” average outcome. All hypothesis tests conducted using a *t*-distribution with degrees of freedom given by the sample size minus the total number of parameters in the propensity score and regression models for each ATT. “Median split” results were calculated by splitting each original covariate at its median, re-estimating the ATT on each subsample, and generating a new estimate of the ATT using a weighted average of the subsample estimates. The median split results are the mean and standard deviation of the resulting ATT estimates. For each ATT, the *q*-value gives the proportion of rejected null hypotheses that would in fact be true if we rejected all null hypotheses with *p*-values no greater than what is reported for the given ATT. The *q*-values were calculated using the method of Benjamini and Hochberg (1995).

Table 5. Difference-in-differences estimates for outcomes observed pre-intervention

	CIA ATT	DID ATT	Percent impact	Standard error	<i>p</i> -value	<i>q</i> -value
<i>Year one</i>						
Irrigated plantain area	0.42	1.26	251.22%	0.69	0.070	0.070
Area planted with plantain	0.07	0.52	67.79%	0.25	0.037	0.074
Area in maize	0.07	0.13	13.63%	0.15	0.397	0.397
Total credit	-1.02	-1.09	-66.26%	1.56	0.486	0.486
<i>Year two</i>						
Irrigated plantain area	0.54	1.27	255.30%	0.48	0.010	0.010
Area planted with plantain	0.11	0.20	21.78%	0.15	0.195	0.390
Area in maize	-0.21	-0.16	-14.52%	0.17	0.357	0.357
Total credit	-0.93	-2.18	-88.72%	0.94	0.022	0.022
Observations	146					

Notes: The "CIA ATT" column gives the estimates of the Average Treatment Effect on the Treated obtained under the Conditional Independence Assumption. Difference-in-difference (DID) estimates are only reported for outcomes observed prior to the start of the RBD program. The set of covariates used in estimating the ATTs by DID does not include lagged outcomes, but is otherwise identical to what was used to obtain the main results shown in tables 3 and 4.

Table 6. Pre-intervention placebo effects

	Placebo ATT	Percent impact	Standard error	<i>p</i> -value	<i>q</i> -value
Irrigated plantain area	0.06	5.96%	0.13	0.664	0.664
Area planted with plantain	-0.06	-5.60%	0.10	0.571	0.571
Area in maize	0.18	19.69%	0.14	0.192	0.384
Total credit	-0.36	-30.02%	0.57	0.535	0.535
Observations	146				

Notes: Placebo estimates of the Average Treatment Effect on the Treated are only estimated for outcomes observed prior to the start of the RBD program. The set of covariates used in estimating the placebo ATTs does not include any outcomes, but is otherwise identical to what was used to obtain the main results shown in tables 3 and 4.