

Millennium Challenge Account-Nicaragua Rural Business Program: Summary of Baseline data for Rice and Plantain Farmers, and Analysis of First Year Program Impacts on Participating Rice Farmers

Conner Mullally
November 17, 2010

Summary:

The Rural Business Development (RBD) program of the Millennium Challenge Account in Nicaragua offers technical training, improved technology, and liquidity to small farmers with the goal of transforming them into commercially viable producers. This paper focuses on two crops included in the RBD program: rice and plantain. I summarize the different aspects of the program, and present descriptive statistics characterizing participant and non-participant households growing each crop. I then turn to estimating the impacts of the first year of the RBD program for rice; plantain farmers had yet to experience any impacts when the data were collected, since trees planted as part of the program had yet to begin producing. Results show that at best, the program had no effect on participants, and at worst, it led to a reduction in rice yields as large as 15 percent relative to the average level of output per unit of land in the sample. Tests of possible reasons for this surprising result do not uncover any patterns. These results should be considered in the context of the drought experienced in the study region in 2009, as rice farmers in the sample do not have irrigated land. If climatic conditions return to normal in 2010, and if the components of the RBD program are better suited a more typical production environment, then we might expect to see positive impacts in the analysis of the second round of data to be collected in 2011.

1. Introduction

In July 2005, the Millennium Challenge Corporation signed a five-year, \$175 million compact with the Government of Nicaragua to support those living in the León and Chinandega region by significantly increasing incomes of rural farmers and entrepreneurs. This document describes the characteristics and impacts on participants of on one component of these efforts: the Rural Businesses Development Project (RBD), a program consisting of liquidity access, technology transfer, and agricultural extension services delivered to small farmers in León and Chinandega. The program is administered by Millennium Challenge Account-Nicaragua (MCA). While the RBD includes a wide variety of agricultural activities, the focus here is on rice and plantain farmers.

In what follows, I summarize the details of the RBD program as they pertain to these two crops. I compare and contrast participants in the program with nonparticipants, using household survey data collected in early 2010. I then analyze the impacts of the RBD program on rice farmers in year one of the program; plantain farmers were in the early stages of the program when the data were collected, and impacts on that group of households will be estimated following the second round of data collection, scheduled to being in January 2011.

The results of the analysis indicate the program impacts were either non-existent or negative, or ranging from zero impact to as much as a 15 percent reduction in rice yields relative to the average level of productivity in the sample. The sign of the estimated impact is consistent across different model specifications, although statistical significance is more sensitive to the assumptions employed. I explore possible reasons for these surprising results, but the data do not seem to capture systematic differences in input choices or planting decisions that might explain these findings.

These results should be considered in the context of the poor production conditions in León and Chinandega in 2009, where an intense drought affected output levels; no rice farmers in the sample have access to irrigation. While the RDB program does not appear to have helped rice farmers in 2009, it could be that the technical assistance and enhanced access to inputs it offers are better suited to more typical production conditions. If this is true, then we should expect to see positive impacts of the program in 2010, as long as precipitation returns to normal levels.

2. RBD Program characteristics

2.1 Program goals and benefits

For the two crops considered in this report, the goals of the RBD program are similar. The primary goal of the RBD program for plantain growers is to transform small farms into viable commercial producers of plantain by increasing output levels, improving output quality, and enhancing the consistency of production levels and quality throughout the year. Low quality is seen as an obstacle to expanding sales into markets outside of Nicaragua, while inconsistent availability of output can exclude farmers from contracting arrangements where a steady flow of product is needed by buyers. The benefits to individual producers of the program include access to high quality inputs such as improved seed, the delivery and installation of micro-aspersion irrigation systems, and technical training. The latter component includes training in the proper use and maintenance of the newly delivered irrigation equipment, and instruction on how to protect the plant from disease and timing harvests so as to maximize production. Participation in the program lasts two years; once an irrigation system is installed, it takes a year for output to be

realized from the newly irrigated area, and as a result impacts on output are not felt until the second year of the program.

It is also the goal of the RBD program for rice to make participants commercially viable, although since rice is an annual rather than a perennial crop, the way in which this problem is approached differs slightly from that of the plantain program. There are essentially two levels of benefits to the RBD rice program: cooperative-level benefits and individual-level benefits. Rice farmers participating in the RBD program are all members of cooperatives, and cooperatives with members in the program receive bundles of inputs for rice production sufficient for three manzanas per participating farmer from MCA. These inputs are then lent out to participating members; interest rates on these loans vary across cooperatives, as credit contract details are controlled by cooperatives rather than MCA.

While the input packets are meant to spur production in the short term, they are also designed to help each cooperative establish a rotating credit fund that will make liquidity available to farmers at in future years. For each participating cooperative, MCA pays a maximum of 30 percent of the costs associated with the program; the rest is paid for by the cooperative.

At the level of the producer, the RBD program for rice farmers also features benefits in the form of agricultural extension services, focused on tailoring the use of chemical fertilizers to the soil characteristics of each individual farm, as well as more efficient use of agrochemicals meant to control threats to the plant. The costs of this technical assistance are factored into the portion of the total cost of program participation borne by each participating cooperative.

In the second year of the program, participating cooperatives are set to receive drying patios (“tendals”) to be used by their members. The price of rice received by the producer is determined in part by the level of moisture in the final product, and incorporating better access to

drying patios should increase the average price per unit received for these farmers. This particular benefit was not part of the first year of the RBD program.

2.2 Selection of participants

The plantain program has existed since 2007; 60 farmers joined in the program in the first year, 124 began participation in the second, and 336 enrolled in 2009. Discussions with leaders of farmer cooperatives indicated that many of their members felt cautious about the program when it was first introduced, but once they saw the benefits enjoyed by participants, demand for space in the program increased greatly. The decision to apply for the program is taken by individual farmers. For an application to be considered, farmers must satisfy the following criteria:

- Must have at least 2 manzanas¹ with access to a water source for year round irrigation.
- Area of farm must be between 2 and 20 manzanas.
- Must have planted at least 2 manzanas of plantain prior to installation of irrigation system.
- The depth of the well for irrigation water cannot exceed 15 meters.
- The main land parcel for plantain must be property of the beneficiary.
- The main land parcel for plantain must be outside environmentally sensitive areas.
- The beneficiary must be at least 20 years of age.

Farmers wanting to participate in the RBD program must submit a “business plan” to the MCA office, detailing how the producer would benefit from participation, how he or she would finance participation, and evidence that he or she has no access to services similar to those provided by

¹ 1 manzana = 1.72 acres = 0.70 hectares

MCA, along with a detailed budget. MCA finances up to 30 percent of the cost of the activities in each business plan. Business plans submitted by farmers meeting the criteria listed above are considered by MCA, conditional on verification of details of the plan in the field.

For rice farmers, participation in the RBD consists of several stages, the first of which is satisfying eligibility for participation in the program. Eligibility criteria for rice farmers are summarized below:

- The producer has planted or currently has at least 2 manzanas of rice.
- Area of farm must be between 2 and 50 manzanas, non-irrigated.
- The main rice parcel must be property of the beneficiary.
- The main rice parcel must be outside environmentally sensitive areas.
- The beneficiary must be at least 20 years of age.

While these criteria were meant to be strictly enforced, experience during the process of data collection indicates that this was not the case, particularly with respect to tenancy status.

Additional details of this problem can be found in the report filed by FIDEG with MCC summarizing the first round of data collection.²

In any case, rice farmers interested in participating in the RBD program submit requests for assistance to their cooperatives. The cooperatives then organize these requests into a single business plan that is submitted to MCA office in Nicaragua for approval. Note that in the case of rice, the plans business plans themselves are at the cooperative level but are essentially collections of requests made by individual farms to participate in the RBD program. Whether or not an individual farmer participates in the program depends upon the decision made by MCA with regard to the business plan submitted by his or her cooperative.

² The report is entitled “Informe Final: Levantamiento Línea de Base 450 Hogares Rurales de Los Departamentos de León y Chinandega Productores Arroz y Plátano.”

3. Data³

3.1 Sample design

This section is an overview of the data collected from rice and plantain households in 2010 as part of the impact evaluation effort. Two separate samples were designed for this evaluation: one consisting of participants in the RBD program for plantain producers and a comparable group of control producers, and a second sample consisting of participating rice farmers and controls. The plantain sample was constructed using a list of producers 239 who began the program in 2009, and 56 producers who applied to participate in the program but were rejected. These were the only lists of plantain producers that were made available by MCA. On one hand, the fact that both lists consisted of farmers who had applied for the program suggests that we might expect these two groups to be fairly similar; this sort of similarity can be important for the construction of a quality control group when assignment to treatment (e.g., participation in the RBD program) is not random.

On the other hand, the precision of estimated impacts hinges in part on balancing the proportion of treated and control households, with a fifty-fifty split being optimal. This was clearly impossible here. Instead, I randomly chose 150 farmers from the list of treated households, and 50 randomly chosen farmers selected from the list of untreated households to serve as controls. The six remaining untreated households and 40 randomly selected treatment households were assigned to a list of replacements to be used when households could not be interviewed or failed to satisfy criteria for participating in the RBD program.

For the rice evaluation, the sample was drawn from a list of rice producers provided by MCA who were members of the 11 cooperatives originally chosen to participate in the RBD

³ Note that descriptions of the data files generated by the survey are included in Appendix 2.

program, and thought to satisfy the criteria listed in above in section 2.2 for program participation. It was decided that three of these cooperatives were located too remotely, and data collection from these households would be impractical. The remaining eight cooperatives served as the basis of the sample. Farmers from one cooperative (“Amprosor”) were oversampled relative to their counterparts. Through conversations with project managers from Chemonics and MCA, I learned that this cooperative lost its eligibility for the project with little time left in the rice planting season for not meeting the legal requirements needed to sign a contract with MCA. This seemed like a source of variation in participation that could plausibly be unrelated to potential household outcomes with and without participation in the RBD program, and members of Amprosor were oversampled in order to take advantage of this possibility in later econometric work. Farmers from the other seven cooperatives were sampled randomly. All other farmers in the list provided by MCA were added to a list of replacements.

3.2 Data collection issues and solutions

Once data collection began, it was discovered that a substantial proportion of farmers in the sample did not satisfy the criteria for participation in the program as described above in section 2.2. In the case of rice, this was largely due to selected households renting all land, rather than owning their primary parcels as required by the program. In the case of plantain, tenancy was not so much the problem, but other violations of the criteria were common, such as planting less than the required 2 manzanas of plantain. **Table 1** is a reproduction of a table from the report by FIDEG on the data collection process, and it summarizes these issues with the sample:

Table 1. Reasons for eliminating households from the sample.

<u>Reason</u>	<u>Rice</u>		<u>Plantain</u>		<u>Total</u>
	<u>Treated</u>	<u>Untreated</u>	<u>Treated</u>	<u>Untreated</u>	
Member of other interviewed household	29	18	2	2	51
Violates some program criterion	1	11	12	11	35
Does not own land	45	16	6	6	73
Could not be located	4	4	0	1	9
Outside country	4	2	1	2	9
Refused interview	1	1	0	0	2
Name repeated in the sample	4	9	2	2	17
<u>Total</u>	88	61	23	24	196

Source: FIDEG

That a large portion of selected households would not satisfy program criteria is not necessarily a problem from the standpoint of evaluation; for example, if it is found that a large proportion of sampled farmers do not actually own land despite the fact that the RBD program requires that they do so, then estimated effects of the program will describe impacts on a group of farmers whose tenancy status tends to differ from what was expected in the early stages of designing the sample. However, MCA insisted that only farmers satisfying the criteria for participation be included in the sample. Given that MCA and FIDEG are now going to carry out an additional round of data collection, this was a wise choice; conversations with MCA indicated that farmers that were found to not satisfy participation criteria will not be allowed to participate in year two of the program, at least in the case of rice. If we would have ignored program criteria in year one of data collection, a much smaller proportion of households listed as RBD participants in year one would show up as in the program in year two. As a result, it might have become necessary to replace a large portion of year one sampled households in year two of data collection, weakening the panel nature of the data set.

The net result of these discrepancies between the information in the lists provided by MCA and reality in the field was a change in the design of the sample. In the case of plantain, the

list of replacements for the original sample was exhausted, with treatment and control households being eliminated at similar rates as shown in **Table 1**. This was despite the fact that the requirement for sown area in plantain was reduced to a single manzana of land. Rather than provide FIDEG with additional treated households to fill out the sample from the list of producers provided by MCA and Chemonics, I decided that the enumerators would look for other households in the vicinity of the members of the original sample. These new additions to the sample would have to satisfy the RBD program criteria, in order to maintain the quality of the control group; 15 such households were interviewed.

In the case of rice, the list of replacement was also exhausted, and the stratified nature of the original sample (oversampling members of Amprosor) was not reflected in the final database. Enumerators located an additional 27 households in the field in order to reach 300 in total. These new households were screened to satisfy program criteria, other than being members of eligible cooperatives.

3.3 Summary statistics

Table 2 below presents summary statistics for the sub-sample of plantain farmers, focusing on variables related to demographics, wealth, and productivity. The definition of some of the variables in the table is obvious, but others require a bit of explanation. “Agricultural wealth” is the sum of the self-reported value of land and capital used in agricultural production, and “Consumer durables” is defined similarly for household goods. “Expected yields in a normal year” measures what producers expect to harvest in a typical year per manzana of plantain; plantain growers record output in units produced, i.e., number of fruit. “Proportion first quality in a normal year” is a measure of the average quality of a harvest in a typical year.

Table 2. Summary statistics for plantain sub-sample

	<u>Full Sample</u>	<u>Participants</u>	<u>Non- participants</u>	<u>Difference</u>
Gender of household head (Female=1)	0.187 [0.0319]	0.198 [0.0381]	0.159 [0.0591]	0.039 [0.068]
Household size	4.827 [0.167]	4.717 [0.198]	5.091 [0.308]	-0.374 [0.387]
Years of education, household head	5.613 [0.402]	6.123 [0.473]	4.386 [0.735]	1.736* [0.724]
Value of consumer durables	23,537 [3413]	27,796 [4022]	13,275 [6242]	14,521** [5054]
Area of owned land in manzanas	5.250 [0.461]	6.318 [0.525]	2.679 [0.816]	3.638** [0.720]
Agricultural wealth (land, animals, implements value in Nicaraguan Córdobas)	631,245 [53157]	749,607 [60852]	346,101 [94449]	403,507** [82365]
Area with irrigation in 2009	2.252 [0.209]	2.830 [0.234]	0.859 [0.363]	1.971** [0.360]
Proportion first quality in normal year	0.446 [0.0141]	0.441 [0.0168]	0.457 [0.0260]	-0.016 [0.031]
Expected yields in a normal year (Units/mz.)	54,267 [1707]	53,900 [2037]	55,153 [3161]	-1,254 [4787]
Years experience plantain	4.073 [0.253]	4.160 [0.302]	3.864 [0.469]	0.297 [0.567]
Mz. of plantain in production	1.532 [0.178]	1.762 [0.210]	0.977 [0.325]	0.785* [0.324]
Proportion with formal loan	0.540 [0.0408]	0.557 [0.0487]	0.500 [0.0755]	0.057 [0.090]
Proportion with informal loan	0.0867 [0.0230]	0.0849 [0.0275]	0.0909 [0.0427]	-0.006 [0.051]
Observations	150	106	43	150

Standard errors in brackets
** p<0.01, * p<0.05

Plantain comes in three qualities, with first quality being the most valuable. “Years experience plantain” summarizes in how many of the last 10 years farmers produced plantain. “Mz. of plantain in production” measures the area with plantain trees that are currently producing; plantain trees need approximately a year before bearing fruit that can harvested, and will eventually age to the point of being unproductive if allowed to do so.

By examining the summary statistics in Table 2, a few key differences between households participating in the RBD program and non-participants can be gleaned. The

household heads of participant household are better educated on average by nearly 2 full years. Participant households appear to be substantially wealthier than their non-participant counterparts; the difference in agricultural wealth between the two groups of 403,507 Nicaraguan Córdobas is approximately \$18,500, a large amount in a country with a per capita income of \$2,800 in 2009. Consumer durables also exhibit a significant difference between the two groups. These measures are likely subject to a large amount of measurement error, but they do suggest that the two groups are quite different along a key dimension.

Area with irrigation is significantly larger in the participating group; this is to be expected, since installation of irrigation technology is a key part of the RBD program for plantain. Productivity indicators do not show any significant differences; in year one of the RBD program, impacts on plantain productivity would not yet have been felt. Finally, we see that the two groups of households do not differ in their access to credit. Here, I have defined a formal loan as a loan from a bank or microfinance entity; this is in contrast to informal lenders, who consist of moneylenders, farmer cooperatives, agricultural input providers, and other businesses.

Note that these are summary statistics are not weighted. The goal here is to create two subsamples, one for participants in the RBD program, and another of non-participants, rather than a single representative sample of plantain farmers in León and Chinandega. As a result, the averages given under the column “Full sample” should not be interpreted as estimates of the means in the entire population, since the composition of the sample does not likely reflect the proportions of participants and non-participants in the overall population.

Table 3 presents summary statistics for the subsample of rice producers, again focusing on demographics, wealth, and agricultural productivity.

Table 3 . Summary statistics for rice sub-sample

	<u>Full Sample</u>	<u>Participants</u>	<u>Non- participants</u>	<u>Difference</u>
Gender of household head (Female=1)	0.08 [0.016]	0.09 [0.022]	0.07 [0.022]	0.021 [0.032]
Household size	5.16 [0.117]	5.07 [0.170]	5.24 [0.162]	-0.172 [0.234]
Years of education, household head	4.39 [0.247]	4.98 [0.356]	3.86 [0.339]	1.12* [0.493]
Value of consumer durables	20,005 [1610]	20,852 [2335]	19,233 [2229]	1,619 [3187]
Area owned land (mz.), start of 2009	14.1 [1.438]	14.73 [2.085]	13.53 [1.990]	1.205 [2.956]
Agricultural wealth (land, animals, implements value in Nicaraguan Córdobas)	1,053,660 [98004]	1,167,836 [141894]	949,666 [135419]	208,892 [195497]
Rice planted area (mz.), 2008	4.92 [0.575]	4.97 [0.757]	4.86 [0.888]	0.114 [1.257]
Planted rice in 2008	0.85 [0.023]	0.846 [0.030]	0.856 [0.035]	0.114 [1.257]
^a Expected yield in a normal year (QQ/mz.)	64.16 [0.976]	66.12 [1.271]	61.46 [1.490]	4.66* [2.027]
^a Experience with rice	6.74 [0.216]	6.64 [0.284]	6.87 [0.333]	-0.222 [0.430]
^b Yields in 2008 (QQ/mz.)	60.94 [1.632]	65.83 [2.090]	54.28 [2.440]	11.55** [3.215]
Proportion with formal loan in 2008	0.46 [0.029]	0.476 [0.042]	0.446 [0.040]	0.03 [0.058]
Proportion with informal loan in 2008	0.277 [0.026]	0.413 [0.036]	0.153 [0.034]	0.260** [0.050]
Observations	300	143	157	300

Standard errors in brackets
** p<0.01, * p<0.05
^a Estimate is for the 247 households (143 participants in RBD) that planted rice in 2009.
^b Estimate is for the 205 households (121 participants in RBD) that planted rice in both 2008 and 2009.

Where possible, variables have been adjusted to reflect quantities before planting in 2009, rather than at the time the household was surveyed, in order to compare treatment and control households prior to any program impacts. For example, in the case of land, the total area owned at the time of the interview was recorded, along with any land that had been added to that total in the past year, or land that had been sold or lost in the same period of time. The variable “Area

owned land (mz.)” listed in Table 3 is the total recorded at the moment of the interview minus net land acquisitions in the prior year.

The differences between participants and non-participant households are not quite so dramatic in the case of rice relative to that of plantain. Heads of participating households are more educated than their non-participant counterparts, as was the case with plantain. Wealth indicators show that the two groups are quite similar along that particular dimension, whether this is measured using consumer durables, land holdings, or the value of all agricultural wealth. Participating households do appear to be more productive, with significantly higher expected yields in a typical year, and substantially higher yields in 2008; however, it should be noted that the yields figure for 2008 only reflects differences for households that planted that year. Productivity figures are given in quintals of output per sown manzana of land. Quintals are the preferred measure of output among rice farmers, and are equal to 46 kg.

Lastly, we see that participating households had much greater access to informal loans in 2008. The data indicate that prior to the RBD intervention participating households were already receiving a large number of loans from their cooperatives. It is unsurprising that MCA would pick cooperatives who have demonstrated the capacity to manage credit to participate in the program, but it also points out that forming revolving credit funds as a component of the program will not introduce an entirely new service to members of those cooperatives. Pre-intervention access to liquidity among these farmers appears to be fairly high.

4. Setup for model of impacts

4.1 *The population and outcome of interest*

I next turn to estimation of the impacts of year one of the RBD program for rice farmers. Impacts on plantain farmers are not estimated, as trees on land that was newly irrigated at the time of data collection had not yet started producing, and the change in output caused by adopting irrigation technology should generate nearly all impacts of the program.

In this evaluation, I focus on the impact of the RBD program on two groups: rice farmers that satisfy the criteria for participating in the RBD program, and rice farmers that actually participated in the program; the latter group is a subset of the former. Note that estimated impacts will capture effects on farmers who planted rice in 2009, rather than the entire population of farmers who meet program criteria; there are 243 such farmers in the sample of 300.⁴

The outcome I focus on is rice yields, or quintals of output per sown manzana of land. I use yields for two reasons: firstly, one of the main goals of the program is to raise productivity, and secondly, because the survey data contain a measure of yields pre and post-RBD. Information on the outcome of interest prior to the beginning of a program is an important tool for assessing the quality of a statistical model when assignment to treatment is not random. Using recall data on other outcomes of interest such as consumption would not be very fruitful, as households are not going to remember expenditures going back a year except for very large items. In contrast, households whose primary activity is agriculture are likely to remember how much of a major crop they produced one year ago, a point that I verified at least anecdotally in

⁴ Four farmers that were not members of eligible cooperatives reported being participants in the RBD program. Their names were cross-checked against databases maintained by MCA in Nicaragua, and this could not be verified. These households were dropped from the sample used in the analysis, leaving 243 rice planters. The results reported are robust to their inclusion, however.

the field through conversations with farmers, representatives of farmer cooperatives, and agronomists working with farmers in the RBD program.

4.2 Estimation strategy

The evaluation of programs where participation is not random is complicated by the fact that outcomes of interest may be correlated with household characteristics which are also driving the participation decision. Suppose we would like to measure the average effect of participating in the RBD program on crop yields for households that planted rice in 2009 and meet program criteria; this is the “Average Treatment Effect” (ATE) of the program for this subpopulation. Merely comparing participants and non-participants will yield a biased and inconsistent estimate of the ATE. For example, if more talented farmers have a higher probability of participating in RBD, and they also have higher crop yields, then a comparison of participant and non-participant households would attribute too great an effect to the RBD program; part of the observed difference in yields ought to be attributed to the difference in farming ability in the two groups. We also might observe the reverse of this situation, as farmers with particularly low productivity might be more inclined to seek out assistance. In this case, we would underestimate the impacts of the RBD program.

Here I will attempt to control for these confounding factors via the Inverse Propensity Score-Weighted Least Squares method (IPS-WLS). This method involves constructing a weight for each treated observation based on the inverse of its estimated probability of participating in the RBD program, and weights for each untreated observation based on the inverse of its probability of not participating. These weights are then used in a weighted least squares

regression. A detailed explanation of the method is given in Appendix 1, but here I offer a brief explanation of its mechanics.

Suppose we want to estimate the ATE. Let d be a dummy variable taking a value of 1 if a household participates in the RBD program, and zero otherwise. Let X represent the vector of observed characteristics in the data. For a household with observed characteristics $X=x$, its probability of participation can be written as:

$$P(d = 1 | X = x) = p(x) \tag{1}$$

This parameter is known in the program evaluation literature as the “propensity score.” It is estimated for each household using the available data and a model appropriate for a binary dependent variable. The estimated propensity score $\hat{p}(x)$ is then used to construct weights for each household. The weight \hat{w}_i for each treated observation is set equal to:

$$\frac{1}{\hat{p}(x)} \tag{2}$$

And for an untreated observation, \hat{w}_i is:

$$\frac{1}{1 - \hat{p}(x)} \tag{3}$$

These are then used in a weighted least squares regression of yields on explanatory variables, including participation in the RBD.

Suppose we have a regression model of yields in 2009, for example:

$$y = \alpha_0 + d\alpha_1 + x'\alpha_2 + d(X - \mu)'\alpha_3 + u \tag{4}$$

where d and X are defined as before, y represents yields in 2009, μ is a vector of means of the X variables, and u is a disturbance term assumed to be uncorrelated with X and d . In this model, the

average direct effect of the RBD program is given by α_1 , while the parameter vector α_3 will capture differences in program impacts due to household characteristics.

If μ contains the means of the variables in X for the population as a whole, then α_1 is the ATE. To see this, take the difference of the expected value of (4) over the whole population when $d=1$ and the expected value of (4) over the whole population when $d=0$; this is the average outcome when everyone receives treatment minus the average outcome when no one does, or the ATE. The $d(X - \mu)' \alpha_3$ term drops out in expectation, leaving α_1 . The ATT is computed similarly by taking expectations over the subpopulation of participants in the RBD program and replacing μ with means for treated households.

OLS would estimate the parameters of (4) by minimizing the sum of squared residuals. In this paper, these parameters will be estimated by minimizing the weighted sum of squared residuals:

$$\text{Min} \sum_{i=1}^N \hat{w}_i \hat{u}_i^2 \tag{5}$$

where \hat{u}_i is the estimated residual.

The main advantage of this approach comes from the “double robustness” property of inverse propensity score weighting when combined with certain types of models, among them linear regression. This property states that if either the model for the propensity score or the linear regression model for the outcome is correct, then inverse propensity score weighting when combined with linear regression will yield consistent estimates of the ATE and ATT (Wooldridge, 2007). The basic explanation for this property is that if the regression model is correct, then weighting will not affect the consistency of the estimated parameters, and if the propensity score model is correct, the mechanics of the formula for the weighted least squares

estimate of α_1 along with the correctly specified weights will yield a consistent estimate of the ATE or ATT, regardless of whether the regression model is properly specified.

4.3 Propensity score fundamentals

While IPS-WLS does have properties that might make it preferable to either method in isolation, it still relies on the assumption that factors that may be driving both the decision to participate and yields can be controlled using observed household characteristics. In the example of differences in farmer ability given above, we could attempt to measure these differences by collecting data on education, past yield outcomes, or what farmers expect to produce in a given year.

To state this assumption more formally, let $y(1)$ represent yields when treatment is received and $y(0)$ denoting yields without participation in the RBD program; note that these are potential outcomes, and what is actually observed will depend on the treatment status of a given household. I assume that:

$$y(1), y(0) \perp\!\!\!\perp d \mid X = x \tag{6}$$

where d and X are as defined before. This is known as the “unconfoundedness assumption,” and it states that potential outcomes are independent of participation in the RBD program conditional on holding the vector X fixed (Imbens, 2004). For example, suppose there is a group of households wherein each unit has the same observed characteristics. The assumption given in equation (6) rules out the possibility of a larger average $y(1)$ for the households in this group that actually end up participating in the RBD as compared to their untreated counterparts in the same group. Since the outcome and participation only depend on X , we can generate consistent

estimates of the ATE by correcting for differences in the distribution of X between the two groups prior to comparing their average outcomes.

If X includes a large number of variables, holding it fixed at all possible values will be extremely cumbersome. Fortunately, we can condition on the propensity score rather than on the entire X vector (Rosenbaum and Rubin, 1983). That is:

$$y(1), y(0) \perp\!\!\!\perp d \mid p(x) \tag{7}$$

In a randomized experiment, the propensity score can be chosen by the researcher; e.g., it could be set to 0.5 for all households by randomly assigning half to treatment and half to control. In the present context where treatment assignment is not random, this probability has to be estimated using an appropriate model.

4.4 The propensity score in the RBD program context

In the case of the RBD program for rice, the treatment decision is somewhat complicated. The sample used in the analysis consists of farmers who planted rice in 2009 and meet the descriptive criteria for eligibility in the program. Among this group are farmers who are ineligible for the program by virtue of not being members of participating cooperatives. Therefore whether a household participates in the RBD is the product of two outcomes: becoming a member of an eligible cooperative, and joining the RBD program conditional on membership in an eligible cooperative.

Let us represent the decision to join an eligible cooperative using the dummy variable *eligible* and acceptance in the RBD program using a dummy variable *RBD*. Planting rice in 2009 will be represented by the variable *rice*. The new expression for the treatment indicator conditional on planting in 2009 is:

$$(d | rice = 1) = (eligible * RBD | rice = 1) \quad (8)$$

For a household with observed characteristics x , the probability of participation is:

$$\begin{aligned} P(d = 1 | X = x, rice = 1) &= P(eligible = 1, RBD = 1 | X = x, rice = 1) = \\ &P(eligible = 1 | X = x, rice = 1)P(RBD = 1 | X = x, rice = 1, eligible = 1) = p(x) \end{aligned} \quad (9)$$

Given this expression, the propensity score is the product of two probabilities that can be estimated from the dataset.

Alternatively, I could have written:

$$\begin{aligned} P(d = 1 | X = x, rice = 1) &= P(eligible = 1, RBD = 1 | X = x, rice = 1) = \\ &P(eligible = 1 | X = x, rice = 1, RBD = 1)P(RBD = 1 | X = x, rice = 1) = \\ &P(RBD = 1 | X = x, rice = 1) \end{aligned} \quad (10)$$

This second specification may appear preferable, as estimating (10) would only require a single step. However, the choice between participating and not participating is only observed for members of eligible cooperatives. If I were to estimate (10) using observed data, the results would be an estimate of $P(RBD = 1 | X = x, rice = 1, eligible = 1)$, which is not the same as the probability given in (10).

As a result, my model of selection into the RBD program is based on (9), and this requires modeling the probability of being a member of an eligible cooperative, as well as the probability of participating in the RBD program conditional on belonging to an eligible cooperative. Both of these probabilities are estimated using a logit model. A logit regression will specify that the probability of a positive outcome (e.g., membership in an eligible cooperative) conditional on observed characteristics $X=x$ is equal to:

$$\frac{\exp(\pi_0 + x' \pi_1)}{1 + \exp(\pi_0 + x' \pi_1)} \quad (11)$$

where π_0 and π_1 are to be estimated. Two such regressions must be estimated in this case, one for the outcome *eligible* and another for *RBD*; the propensity score for each household will be the product of the predicted probabilities generated by the two logit regressions.

There is very little guidance in the literature as to how one should go about selecting the proper specification for the propensity score; see Caliendo and Kopeinig (2005) for a review. The most important criterion is that the propensity score model includes all variables that are correlated with the decision to participate in a program and the outcome of interest. While this does not rule out including variables that affect participation but not the outcome, doing so can sharply increase variability of estimated impacts without any gain in terms of reduced bias (Brookhart et al., 2006). Therefore I focus on specifying a good model of yields in 2009, and use the variables included in that model as the basis of the logit models for *eligible* and *RBD*.

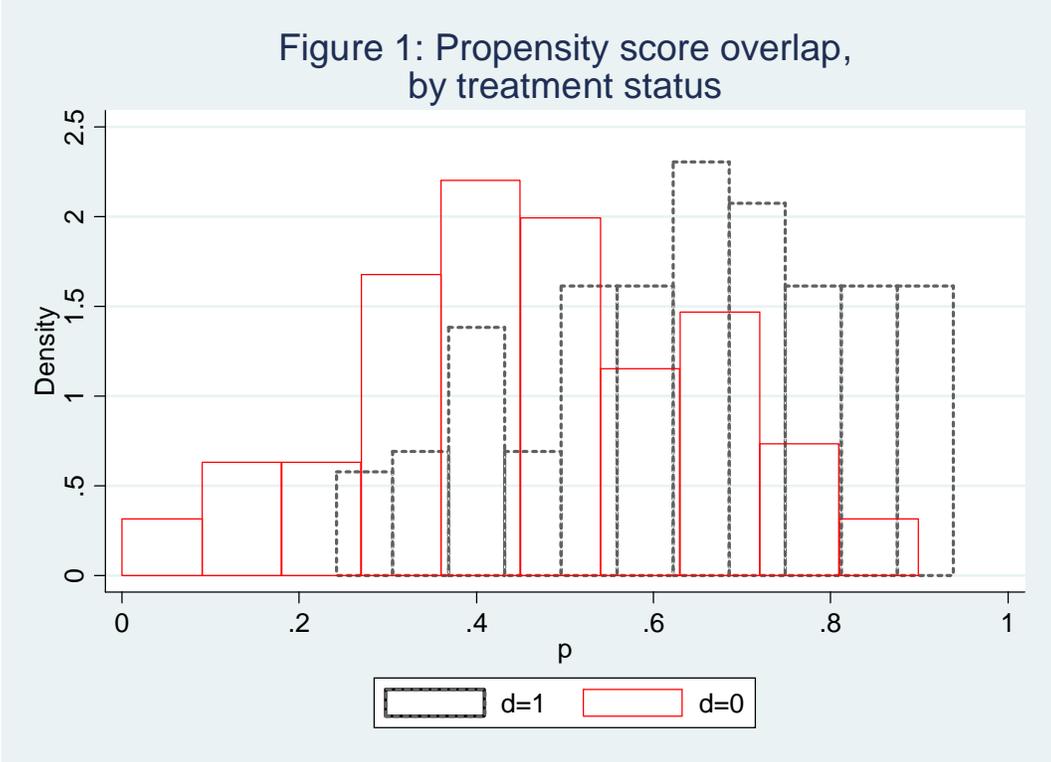
4.5 Model structure: regression

The structure of the linear regression model used here follows the example given in (4). This structure may appear restrictive, but interactions and polynomial terms can be incorporated as elements of the vector X if desired. Given this basic setup, all that remains is selection of the variables to be contained in X . As a first stage in variable selection, I estimated a regression of the form given in (4) including all variables that might plausibly be correlated with y and d while not being impacted by d . Variables whose effects on y in this regression were estimated very imprecisely (i.e., a t-statistic of less than 1) were discarded, unless there was a compelling theoretical reason to keep them in the model. Quadratic terms of the remaining variables were added unless logic precluded their inclusion and kept if they increased the explanatory power of the model; nearly all such higher order terms were discarded. Variables surviving these

specification rules entered into the vector X for estimation of (4), as well as the logit regressions for *eligible* and *RBD*.

5. Results of propensity score estimation

The key diagnostics in assessing an estimated propensity score are the degree of overlap in the estimates across the two treatment groups, and the balance observed in the covariates once they have been corrected for differences in the propensity score. Figure 1 is a histogram depicting estimated propensity scores for treated and untreated observations in the sample.



As described in section 4.4 , these estimates are the product of predicted probabilities generated by separate logit regressions; the logit results are not reported here, but are available upon request. The degree of overlap in Figure 1 is fairly strong, although the treated propensity scores have more density in the right hand tail while the opposite is true of the untreated observations,

which is unsurprising. Overall, the visual impression from the figure is the two groups appear similar enough that any differences can be corrected via controlling for observed characteristics, and that a good model of the propensity score should succeed in balancing the distribution of variables thought to be correlated with d and yields.

Table 4 presents summary statistics, both with and without adjusting the covariates by weighting observations according to the formulas in (2) and (3).

	<u>Raw Averages, by treatment</u>		<u>Raw</u>	<u>Weighted</u>
	<u>status</u>		<u>Difference</u>	<u>Difference</u>
	<u>d=1</u>	<u>d=0</u>		
Planted rice, 2008 (Yes=1)	0.832	0.673	0.159	0.001
	[0.037]	[0.034]	[0.050]***	[0.055]
Had a formal loan, 2008 (Yes=1)	0.46	0.472	-0.012	-0.028
	[0.043]	[0.049]	[0.065]	[0.073]
Had an informal loan, 2008 (Yes=1)	0.416	0.17	0.246***	0.024
	[0.038]	[0.043]	[0.058]	[0.072]
Total credit, 2008 (Córdobas)	33,530.29	102,703.11	-69,172.82	-34,025.73
	[38,834.634]	[44,149.579]	[58,798.930]	[31,182.322]
Altitude of main parcel (m)	106.62	156.179	-49.559***	-9.374
	[7.790]	[8.857]	[11.795]	[12.292]
Altitude*Altitude	16,780.52	36,300.31	-19,519.793***	-4,542.51
	[3,349.309]	[3,807.699]	[5,071.138]	[4,559.128]
Distance from center of León (km)	56.195	44.681	11.514***	0.443
	[2.421]	[2.753]	[3.666]	[4.361]
Gender of household head (Female=1)	0.088	0.094	-0.007	0.006
	[0.025]	[0.028]	[0.037]	[0.037]
Age of household head	47.27	48.028	-0.758	-0.45
	[1.093]	[1.242]	[1.654]	[1.758]
Household size, minus dependents	3.212	3.075	0.136	0.019
	[0.120]	[0.136]	[0.181]	[0.192]
Proportion of rice area affected by shocks, 2009	0.832	0.721	0.111**	0.006
	[0.033]	[0.038]	[0.051]	[0.055]
Expected yield in a normal year (QQ/mz)	66.095	61.575	4.519**	1.071
	[1.308]	[1.487]	[1.980]	[2.041]
Expected yield in a bad year (QQ/mz)	38.526	37.236	1.29	0.178
	[1.450]	[1.649]	[2.196]	[2.154]
Index of agricultural wealth	0.06	0.24	-0.18	-0.128
	[0.168]	[0.191]	[0.255]	[0.231]
Index of non-agricultural wealth	0.276	0.018	0.258	0.083
	[0.190]	[0.216]	[0.288]	[0.329]

Standard errors in brackets
 *** p<0.01, ** p<0.05, * p<0.1

After weighting, covariates are much more balanced across the treated and untreated households. Nearly all raw differences are reduced in absolute value by the IPW procedure, and no significant differences remain. Note that these variables serve as explanatory factors in most of the regressions that follow. Many of these variables are from 2008 rather than 2009, because participation in the RBD program would have impacted observed values in 2009. For example, the variable “Had an informal loan, 2008” would bias estimated program impacts if it contained loans from 2009 rather than 2008, since access to credit through farmer cooperatives (which are informal lenders) is an important component of the RBD program.

The wealth indicators “Index of agricultural wealth” and “Index of non-agricultural wealth” are based on wealth levels prior to program impacts. This required reconstructing asset ownership prior to participation in the RBD. This was done via survey questions that asked households about the number of assets owned at the time of the interview, as well as purchases and losses in the past year. Very minor assets for which we can expect large changes were dropped. For each type of asset, the quantity owned in 2008 was recovered as the number of items owned at the time of the interview, minus net additions to that number in the prior year. Dozens of assets and wealth indicators were recorded, making some sort of aggregation necessary for inclusion in the model. Self-reported asset values could not be used, as these reflect values at the time of the interview. Instead, the indices were constructed using Principal Components Analysis (;). The PCA transformation mapped wealth variables into a series of orthogonal components, each capturing a successively smaller share of the variation in wealth throughout the sample. The indices used here are comprised of the first principal components from separate PCA transformations of indicators of agricultural wealth (including land and agricultural implements) and non-agricultural wealth (housing characteristics, access to utilities

and sanitation, and consumer durables). These indices explain 26.23 percent and 30.25 percent of variation in agricultural and non-agricultural wealth in the sample, respectively.⁵

Note that among the productivity indicators, only variables capturing expected output in normal and bad years are included; expected output in good years was very weakly correlated with yields as well as the participation decision, and was dropped from the model. This is likely due to the fact that León and Chinandega suffered a drought in 2009 that according to MCA was extremely harmful to production. Productivity in relatively good years may be unimportant when production conditions are so poor.

6. Regression results

6.1 Baseline specifications

Table 5 reports regression results for models taking the form given in (4). I estimated weighted least squares (WLS) models as well as weighted Tobit models, and results were computed by bootstrapping the two-step estimation process of estimating both the propensity score and the yields model given in (4). The bootstrapped standard errors listed in Table 5 account for the fact that using weights based on the estimated propensity score may increase the variability of the estimated regression coefficients. Note that these bootstraps resample individual observations; a cluster bootstrap based on cooperative membership yielded more precise results, and I use the individual-level bootstrap to keep inferences conservative.

The Tobit results are reported for the sake of thoroughness. Yields are a censored outcome cut off at zero, and this could be problematic for the WLS model if it were to predict negative outcomes for a non-negligible portion of the sample. Fortunately, this is not the case, and the Tobit results largely reflect those of the WLS estimates.

⁵ Hardle (2007) offers a more detailed explanation of PCA with examples of applications.

Table 5. Regression results

Dependent variable: Rice yields, 2009	(1)	(2)	(3)	(4)
	<u>ATE, WLS</u>	<u>ATT, WLS</u>	<u>ATE, Tobit</u>	<u>ATT, Tobit</u>
d (Participants in RBD=1)	-5.067*	-4.976	-5.178*	-4.952
	[2.700]	[3.044]	[2.900]	[3.276]
Planted rice, 2008 (Yes=1)	6.118	5.794	6.383	6.172
	[4.996]	[5.514]	[5.410]	[6.074]
Had a formal loan, 2008 (Yes=1)	-1.908	-3.719	-1.715	-3.751
	[4.510]	[5.802]	[4.480]	[5.754]
Had an informal loan, 2008 (Yes=1)	4.561	4.274	5.146	4.763
	[4.653]	[6.102]	[4.824]	[6.427]
Altitude of main parcel (m)	0.259***	0.276***	0.269***	0.282***
	[0.065]	[0.084]	[0.065]	[0.086]
Altitude*Altitude	-0.001***	-0.001***	-0.001***	-0.001***
	[0.000]	[0.000]	[0.000]	[0.000]
Distance from center of León (km)	0.12	0.152	0.128	0.157
	[0.083]	[0.097]	[0.084]	[0.098]
Gender of household head (Female=1)	0.716	-3.528	1.592	-2.54
	[6.266]	[7.740]	[6.500]	[8.095]
Age of household head	0.303	0.354	0.329*	0.376
	[0.186]	[0.251]	[0.193]	[0.260]
Household size, minus dependents	-0.297	0.588	-0.247	0.666
	[1.275]	[1.594]	[1.311]	[1.646]
Proportion of rice area affected by shocks, 2009	-16.947***	-20.043***	-16.501***	-19.570***
	[4.828]	[5.791]	[5.019]	[6.077]
Expected yield in a normal year (QQ/mz)	0.492**	0.312	0.486**	0.305
	[0.223]	[0.286]	[0.232]	[0.302]
Expected yield in a bad year (QQ/mz)	0.459**	0.557**	0.518**	0.624**
	[0.224]	[0.273]	[0.238]	[0.294]
Index of agricultural wealth	-1.031	-1.979	-0.89	-1.749
	[1.277]	[1.630]	[1.288]	[1.663]
Index of non-agricultural wealth	-0.147	-0.325	-0.147	-0.385
	[0.953]	[1.136]	[1.037]	[1.273]
<u>Interactions</u>				
d*(Planted rice, 2008)	-4.794	-2.095	-4.85	-2.165
	[6.846]	[7.033]	[7.499]	[7.704]
d*(Had a formal loan, 2008)	3.83	5.867	3.493	5.932
	[5.377]	[6.447]	[5.398]	[6.442]
d*(Had an informal loan, 2008)	0.923	0.887	1.14	1.016
	[5.874]	[6.994]	[6.162]	[7.370]
d*(Altitude of main parcel)	-0.003	-0.551	0.012	-0.006
	[0.094]	[2.086]	[0.099]	[0.113]
d*(Altitude*Altitude)	0	-0.022	0	-0.12
	[0.000]	[0.108]	[0.000]	[0.116]
d*(Distance from center of León)	-0.101	-0.117	-0.107	0
	[0.103]	[0.113]	[0.107]	[0.000]
d*(Gender of household head)	8.278	0	7.853	7.742
	[11.316]	[0.000]	[12.013]	[11.552]

d*(Age of household head)	-0.282 [0.257]	8.457 [10.916]	-0.295 [0.265]	-0.272 [0.319]
d*(Household size, minus dependents)	0.779 [1.897]	-0.266 [0.310]	0.735 [1.961]	-0.636 [2.156]
d*(Proportion of rice area affected by shocks, 2009)	-7.377 [6.476]	-6.476 [6.979]	-7.825 [6.694]	-7 [7.246]
d*(Expected yield in a normal year (QQ/mz))	-0.328 [0.270]	-0.158 [0.328]	-0.273 [0.290]	-0.111 [0.350]
d*(Expected yield in a bad year (QQ/mz))	0.132 [0.255]	0.048 [0.296]	0.062 [0.271]	-0.026 [0.317]
d*(Index of agricultural wealth)	-1.57 [1.686]	-0.303 [1.980]	-1.774 [1.755]	-0.558 [2.039]
d*(Index of non-agricultural wealth)	0.136 [1.267]	0.598 [1.371]	0.189 [1.360]	0.722 [1.502]
Constant	-29.658* [16.042]	-26.358 [19.956]	-35.846** [17.297]	-32.189 [21.349]
Observations	243	243	243	243
R-squared (based on single estimate using full sample)	0.6893	0.699	0.1255	0.1286

Bootstrapped standard errors in brackets, 999 replications
*** p<0.01, ** p<0.05, * p<0.1

The WLS estimate of the ATE given in column (1) is negative and significant at 10 percent (p-value = 0.061). Note that the ATE does not represent estimated impacts on actual participants, but rather what the average impact would have been if all farmers in the sample had participated. The estimated ATE indicates that if all households in the sample were to participate in the RBD program, yields would have decreased by an average of 5.1 quintals per sown manzana. Average yields in the sample in 2009 were 48.8 quintals/manzana, so this figure represents approximately a 10 percent decrease from the average level of productivity.

The WLS estimate of the ATT, which capture average impacts on participants, is also negative but is not quite precise enough to pass for significant at acceptable levels (p-value= 0.102). It is slightly higher but nearly equal to the ATE. Since the ATT is estimated somewhat less precisely than the ATE, I cannot state unequivocally that based on these results, the RBD program reduced yields in 2009. But these regression results do constitute evidence for this

conclusion, and certainly do not support the notion that the RBD program raised productivity in 2009.

Most of the covariate effects are estimated imprecisely, with a few exceptions. Having a main rice land parcel at a higher altitude is associated with higher yields; when evaluated at the sample average altitude of 128 meters above sea level, an extra meter of altitude results in 0.13 quintals of additional production per sown manzana. The proportion of sown area affected by exogenous shocks (primarily drought) had a strong and negative effect on yields. Note that this variable entered into the regression as a proportion, and the proper way to read the estimated effect of shocks on yields in Column (1) is that a 1 percent increase in the share of planted rice area affected by shocks was associated with a decrease in yields of 0.169 quintals per manzana. The estimated coefficients on the interactions between participation in the RBD and observed characteristics are all insignificant. These interactions are meant to capture how the impact of the treatment varies according to observed characteristics; based on these results, there is no evidence that program impacts on yields varied with household characteristics in any discernible way.

The signs and magnitudes of the estimated ATE and ATT parameters are surprising to say the least. Given the impact of the drought in 2009, we might expect to observe no impact of the RBD program, but a negative impact is counter intuitive. In what follows, I subject these results to a series of robustness checks in order to detect any sensitivity of the results to modeling assumptions. These tests consist of assessing the unconfoundedness assumption, dropping outliers, and eliminating variables from the model that may distort estimated impacts.

6.2 Robustness checks: evidence for the unconfoundedness assumption

In this section, I conduct indirect tests of the unconfoundedness assumption, i.e., potential outcomes are independent of treatment assignment conditional on observed characteristics, and then adjust the model based on these results when necessary. This assumption can never be tested directly, but it is possible to provide evidence for or against it via indirect tests. As stated in section 4.2, a properly specified model for the impacts of the RBD program ought to show no impact of the program on yields in the previous year; farmers were not yet participating, and therefore should not yet have been affected. To examine this assumption, I estimate the same model used to generate columns (1) and (2) of Table 5, but I replace the outcome with yields from 2008, and drop the explanatory variable “Planted rice, 2008.” Unfortunately, this means dropping 49 farmers from the model, as not all households planting in 2009 did so in 2008. The results of the test are still suggestive, however. If they show a program impact on the lagged outcome, it seems unlikely that this effect would disappear if more households were included in the estimating sample.

Additional evidence for or against the unconfoundedness assumption in this case can be generated by regressing yields in 2009 on a treatment that ought not to have affected the outcome. For example, suppose we look at the subsample of non-participating members of eligible cooperatives and ineligible households. Let the indicator variable e take a value of 1 for the former group and 0 for the latter. If the model is correctly specified, the variable e should have no effect on the outcome. To test this, I estimate a model of the same form as what is reported in Table 5, but on the subsample of non-participating households and replacing the treatment indicator with the variable e . The results are presented below in Table 6, omitting covariates and interactions.

Table 6. Assessing unconfoundedness				
Dependent variable:	Yields in 2009		Dependent variable:	Yields in 2008
Estimation sample:	Non-participants		Estimation sample:	Rice planters in 2008 and 2009
Treatment:	Eligibility		Treatment:	RBD program
	(1)	(2)	(3)	(4)
	<u>ATE of eligibility</u>	<u>ATT of eligibility</u>	<u>ATE of RBD, lagged yields</u>	<u>ATT of RBD, lagged yields</u>
	-0.255	0.378	3.775	8.346
	[12.691]	[23.596]	[4.904]	[5.657]
Observations	106	106	204	204

Bootstrapped standard errors in brackets, 999 replications

The results generally support the unconfoundedness assumption. In the case of the model comparing the two groups of non-participants, the low sample size likely means that even if there were differences between the two groups, it would be difficult to detect them due to low power. Regardless, the magnitude of the estimated effect is very small. It is interesting to note that while there is no significant correlation between yields in 2008 and the treatment indicator conditional on observables, there is some evidence of a positive effect in the ATT model (p-value of estimated impact = 0.14). This suggests that among the subpopulation of households that planted in both years, participants had higher yields in 2008. If we interpret this as a sign that farmers participating in the RBD program tend to be more productive than their non-participant counterparts, the estimates reported in Table 5 would be biased upwards.

6.3 Incorporating lagged yields into the model

The results in Table 6 appear to support the inferences made based on Table 5, but just to be safe, I examine how incorporating rice yields from 2008 into the model affects results. The way in which we interpret estimates that incorporate lagged yields hinges on what we assume

about the productivity of participants versus non-participants, and whether any differences in productivity across the two groups are driven by unobserved heterogeneity. In the present context, it appears reasonable to assume that participants are more productive than non-participants. The variables “Expected yields in a normal year” and “Expected yields in a bad year” are both higher for participants on average, as are yields in 2008; this does not rule out unobserved components of productivity that are negatively correlated with participation in the RBD program, but it is suggestive.

There are two ways in which I could incorporate lagged outcomes into the model of RBD program impacts. The first is to use the difference of yields in 2009 and yields in 2008 as the dependent variable in a regression on d , covariates, and interactions with d . The differenced model is appropriate if there are time-invariant, unobserved differences in productivity that are correlated with participation in the RBD program as well as yields; taking the difference between yields in 2008 and 2009 eliminates this potential source of bias. Note that the estimated program impact from this model will still capture the effect on yields in 2009, rather than the impact on yield growth.

The second option is to keep yields in 2009 as the dependent variable, and add yields in 2008 to the list of covariates. The lag model is appropriate if there are no fixed effects, but lagged yields are positively correlated with participation as well as yields in 2009. While there is no way to tell if the differenced or the lag model is correct, estimated treatment effects from the two models will bracket the true average effect of the program if either of the scenarios described above holds (Angrist and Pischke, 2009). Therefore I report the results of each model in Table 7 below, omitting covariates and interaction terms.

Table 7. Incorporating yields in 2008 into the model

	<u>Lag Model</u>		<u>Differenced Model</u>	
	(1)	(2)	(3)	(4)
	<u>ATE</u>	<u>ATT</u>	<u>ATE</u>	<u>ATT</u>
d (Participants in RBD=1)	-6.538**	-6.473	-9.555	-13.651**
	[2.987]	[4.049]	[5.897]	[6.063]
Observations	204	204	204	204
R-squared (based on whole sample)	0.7068	0.7091	0.4884	0.4936

Bootstrapped standard errors in brackets, 999 replications

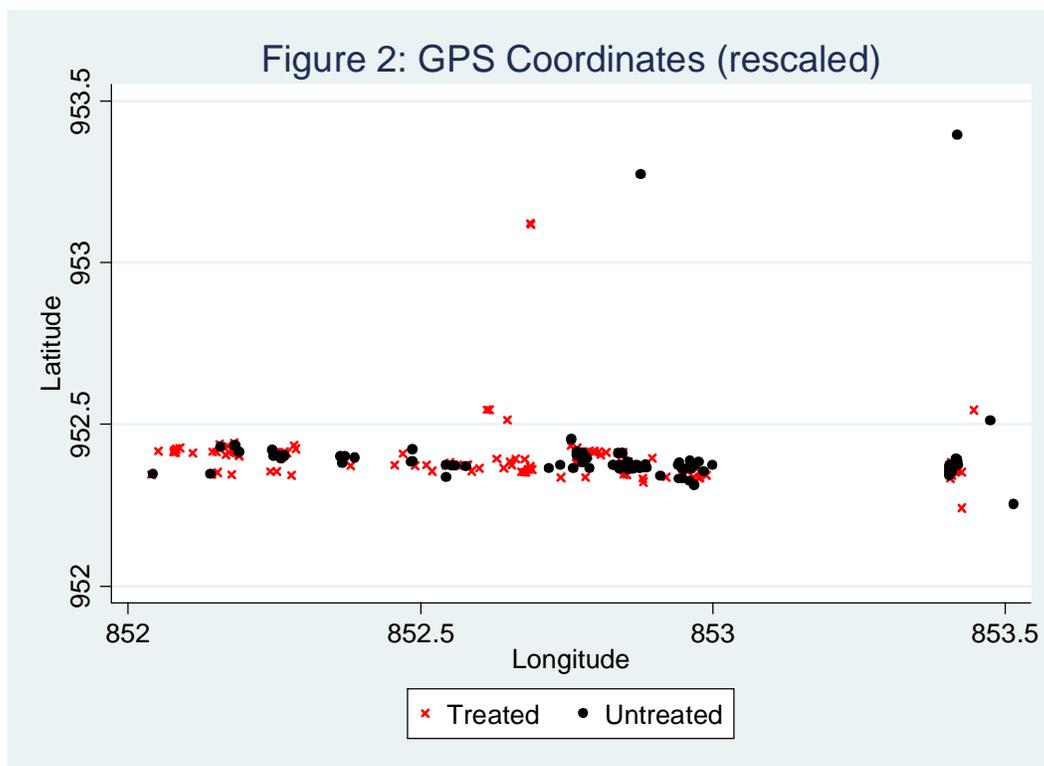
*** p<0.01, ** p<0.05, * p<0.1

Columns (1) and (2) list the estimated ATE and ATT for the model regressing yields in 2009 on yields in 2008 and the full set of covariates and interactions listed in Table 5, with the exception of “Planted rice in 2008.” Columns (3) and (4) report results for the differenced model. Note that when yields in 2008 are included among the covariates, they are also included in the model of the propensity score, but this is not the case when the dependent variable is the change in yields from 2008 to 2009.

If either set assumptions described above are true, then the results in Table 7 bracket the true ATE and ATT for the subpopulation of farmers planting rice in both years. The ATT for the lag model (p-value = 0.11) and the ATE for the differenced model (p-value = 0.105) are estimated imprecisely, so a conservative conclusion based on Table 7 would be that the true average effect of the program is less than or equal to zero, both for the whole population of farmers planting rice in 2008 and 2009 and for the subset of farmers in this same population that participated in the RBD program.

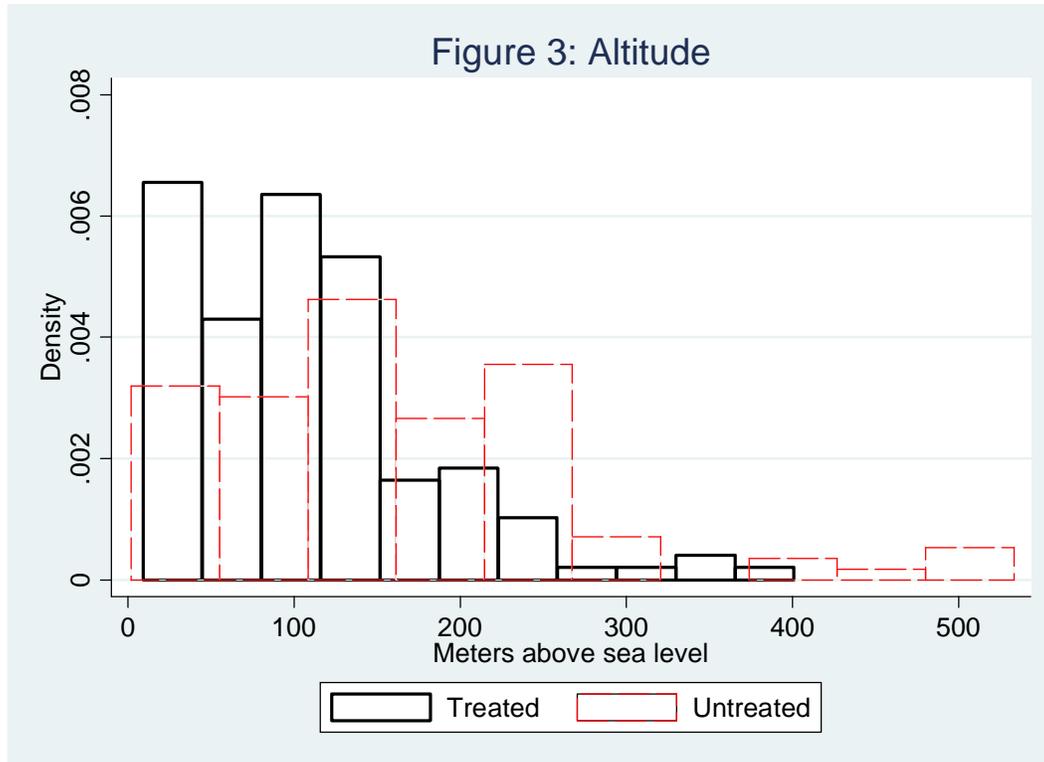
6.3 Robustness checks: geographic outliers

As stated earlier, León and Chinandega suffered a drought in 2009. While I have tried to control for impacts of exogenous shocks, it is possible that the intensity of these shocks varied due to the geographic characteristics of the farms in the sample. There are a handful of geographic outliers that could be discarded as a robustness check. Figure 2 is a crude map showing the rescaled latitude and longitude coordinates of the main rice parcels for households by treatment status.



In general, households are quite close together, but a small number clearly do not fit the pattern of high population density.

Figure 3 below is a histogram showing the distribution of altitude by treatment status.



The histogram for untreated households has more density at higher elevations, and the results in Table 5 indicate that yields are an increasing function of altitude. Therefore to control for the possible impact of geographic outliers, I drop households with parcels located above 400 meters above sea level and households above a rescaled latitude of 953, and then estimate the models from Columns (1) and (2) of Table 5. The results are shown in Table 8, with estimated coefficients for covariates and interactions omitted.

Table 8. No geographic outliers		
	(1)	(2)
Dependent variable: Rice yields, 2009	ATE, no geographic outliers	ATT, no geographic outliers
d (Participants in RBD=1)	-5.126*	-4.973*
	[2.753]	[3.020]
Observations	234	234
Bootstrapped standard errors, 999 replications		
*** p<0.01, ** p<0.05, * p<0.1		

The estimated ATE and ATT are robust to the exclusion of geographic outliers. The estimated impacts maintain their signs and magnitudes, and the ATT is now estimated more precisely and is significant at the 10 percent level. This does not completely rule out the possibility that some sort of spatial heterogeneity could be driving participation in the RBD as well as yields; e.g., there could be systematic differences in soil quality across the two treatment groups that fit this description. However, I can only test what is available in the data, and it does not contain additional detail on heterogeneity with respect to production conditions.

6.3 Robustness checks: dropping variables potentially impacted by the RBD program

In the regression specifications above, I have included expected yields in normal and bad years as explanatory variables. It is possible that the RBD program had permanent and positive impacts on farmer productivity, and that when asked these questions, participant farmers included these impacts when giving their answers. If this is the case, then holding these variables constant in the WLS regressions would bias the estimated impacts downwards; the estimates would capture average program benefits, net of permanent increases in productivity.

If the RBD program affected farmer productivity then it should be possible to detect these impacts by applying the IPS-WLS model, where productivity indicators would replace yields as the outcome of interest. Evidence of positive impacts of the RBD on productivity would necessitate re-estimating the model of yields without including productivity indicators, and checking for any changes in the estimated ATE and ATT.

Table 9 below reports results of using “Expected yields in a normal year” and “Expected yields in a normal year” as dependent variables in the IPS-WLS, with Columns (1) and (2) listing results for the former and Columns (3) and (4) for the latter. The full set of covariates and

interactions used in Table 5 was used, minus the productivity indicators that are now dependent variables. While none of these results show significant impacts, there is some evidence that the RBD program may have affected productivity in normal years; the p-value of the estimated ATE of the RBD program on productivity in a normal year is 0.21 in Column (3), whereas for the ATT in Column (4) it is 0.17, while the magnitudes of both estimated effects are large.

Table 9. Impacts on self-reported productivity				
	<u>Expected yields in a bad year</u>		<u>Expected yields in a normal year</u>	
	(1)	(2)	(3)	(4)
	<u>ATE</u>	<u>ATT</u>	<u>ATE</u>	<u>ATT</u>
d (Participants in RBD=1)	0.215	0.801	3.284	4.368
	[2.453]	[2.429]	[2.618]	[3.209]
Observations	243	243	243	243

These results can be interpreted in two ways. The first is that the RBD program had positive effects on productivity in a normal year, and therefore the regression results reported in Table 5 should be checked for sensitivity to the exclusion of “Expected yields in a normal year” from the list of covariates. The alternative interpretation is that the estimates in Table 9 are capturing pre-program differences in productivity, similar to what was found earlier when using yields in 2008 as the dependent variable. If this second scenario is true, then the results in Columns (3) and (4) suffer from omitted variable bias, as will the estimates of impacts of the RBD program on yields if “Expected yields in a normal year” were dropped from the list of covariates. Assuming that productivity is positively correlated with participation in the RBD program, yields in 2009, and “Expected yields in a normal year,” this bias will be positive in all cases.

If Table 9 is capturing pre-program differences in productivity rather than program impacts, then substituting an alternative measure of productivity that could not have been affected by the program for “Expected yields in a normal year” in the model of yields in 2009

ought to closely match replicate the results of estimating the model for yields while including “Expected yields in a normal year.” The data do not contain an adequate alternative measure for the whole sample of rice growers, but they do contain lagged yields for those that planted in both years. A strong indicator that there were no impacts of the RBD program on productivity would be little change in the estimated impacts of the program generated by the differenced and lagged yields models reported in Table 7, after dropping “Expected yields in a normal year” from the list of explanatory variables.

Table 10 reports the results of re-estimating the WLS models used to generate Table 5 (full set of covariates and interactions with no lagged yields) and Table 7 (the lag model and the differenced model) after dropping “Expected yields in a normal year” from the list of explanatory variables.

	<u>Baseline model</u>		<u>Lag model</u>		<u>Differenced model</u>	
	(1)	(2)	(3)	(4)	(5)	(6)
	<u>ATE</u>	<u>ATT</u>	<u>ATE</u>	<u>ATT</u>	<u>ATE</u>	<u>ATT</u>
d (Participants in RBD=1)	-3.88	-3.705	-4.879*	-4.756	-9.351	-12.338**
	[2.599]	[2.808]	[2.917]	[3.327]	[6.147]	[5.364]
Observations	243	243	204	204	204	204
R-squared (based on full sample)	0.6826	0.6972	0.7007	0.7100	0.4682	0.4536

*** p<0.01, ** p<0.05, * p<0.1

The results in Table 10 show an interesting pattern. Dropping “Expected yields in a normal year” from the baseline model as is done in Columns (1) and (2) pushes the estimated ATE and ATT closer to zero, and they are no longer significant. Estimating this same model while including yields in 2008 in the list of covariates has a similar effect on the estimated ATE and ATT; results from this model are reported in Columns (3) and (4). The ATE remains negative and significant at a 10 percent level in the lag model, but the magnitude of the ATE is increased by around 2 quintals per manzana. The results in Columns (1) through (4) are consistent with the possibility

of negative impacts or zero impacts of the RBD program in 2009, but permanent increases in productivity in a typical year. They are also consistent with the possibility of omitted variable bias, i.e., unobserved components of productivity biasing estimates upwards.

The results in Column (5) and (6) are for the differenced model. Whereas the results of the baseline WLS model and the lagged yield model were quite sensitive to dropping “Expected yields in a normal year,” the differenced model hardly budges. The ATE from this model is imprecisely estimated (p-value= 0.128), but the ATT is significant at the 5 percent level. This suggests that after removing time-invariant differences in productivity by taking the difference of yields in the two years, “Expected yields in a normal year” does not control for any additional differences among farmers. Thus it appears that for farmers who planted in both years, the RBD program had no discernible impact on productivity, and in fact decreased yields in 2009. For farmers who only planted in 2009, the impact is less clear, but it seems safe to say that it was not positive.

7. Unpacking the results

7.1 Possible explanations

If the RBD program did not accomplish its goals in terms of increased productivity, it would be helpful to know why. The most obvious factor depressing its impacts is the drought that occurred in 2009, and it could be that if climatic conditions returned to normal in 2010 then the program will generate positive impacts. But the drought does not explain negative impacts of the program. It can always be argued that the models estimated above, which rely on observed household characteristics, are not controlling for time-varying heterogeneity in productivity that is driving down impacts. But this seems unlikely, given that treated households had substantially

higher yields than their untreated counterparts in 2008. It could be that treatment households have higher output when production conditions are good yet suffer more when the climate is not conducive to growing rice. In any case, explanations based on heterogeneity not captured by the data cannot be tested. But there are other possibilities that can be examined given the information that is available.

7.2 Expansion of sown area onto marginal lands

One possibility is that by participating in the RBD program, households were induced to expand their sown area of rice, possibly onto land that is not profitably used in this way in drought conditions. An implication of this hypothesis can be tested by looking at impacts of the RBD program on sown area in rice in 2009. In Table 11 below, I report the results of regressing sown area in 2009 on participation in the RBD program, covariates, and interactions with the RBD program. I use the same models that served as the basis of measuring impacts on yields in 2009: IPS-WLS of the outcome on a full set covariates and interactions, IPS-WLS of the outcome on covariates, interactions, and the lagged outcome, and IPS-WLS of the differenced model, with inference based on the bootstrap.

Table 11. Impacts on sown area						
	<u>Baseline model</u>		<u>Lag model</u>		<u>Differenced model</u>	
	<u>ATE</u>	<u>ATT</u>	<u>ATE</u>	<u>ATT</u>	<u>ATE</u>	<u>ATT</u>
	(1)	(2)	(3)	(4)	(5)	(6)
d (Participants in RBD=1)	2.363 [1.871]	0.713 [1.038]	-0.188 [0.865]	-0.268 [0.878]	-0.459 [0.854]	-0.782 [0.994]
Observations	243	243	243	243	243	243
R-squared (based on full sample)	0.4100	0.4529	0.7269	0.7467	0.406	0.4062
Bootstrapped standard errors in brackets, 999 replications						

The results suggest that expansion of sown area is not the culprit. In general, the magnitudes of the estimated coefficients are quite small, and none comes close to being significant at an acceptable level.

7.3 Composition of input expenditures

Another possible explanation for the lack of positive impact of the RBD program could be found in data on input expenditures. An important component of the technical assistance offered by the RBD program is instruction on the proper use of fertilizer based on soil chemistry. In addition, the RBD program is setup to delivery additional liquidity to participant farmers, which is to be used to purchase the inputs needed to carry out individual business plans. It could be that in times of drought, increasing the use of certain inputs, e.g., fertilizer, is detrimental to production. In any case, a systematic difference in input use across the two treatment groups can be detected could be interpreted as evidence that the changes in agricultural practices caused by the program were not suitable for drought conditions.

It should be noted that the available data on input expenditures are likely subject to a very high degree of measurement error, making any estimated impacts more variable than would be the case with better quality data. The data are based on retrospective questions asking about inputs used in the production of rice in the 2009 growing season; accurate input data would likely require better timing of interviews with multiple return visits throughout the year to record changes in input use. Also, price data were provided by the MCA office in León, and are market prices. But there may be a substantial degree in heterogeneity in the prices actually paid by farmers for inputs. For example, a farmer receiving pesticides as part of a loan may not pay the same price as his neighbor who buys them at an input store.

Table 12 lists the results of regressing different measures of input use intensity on participation in the RBD program, a full set of covariates (not including lagged yields or “Expected yields in a normal year”), and interactions.

Table 12. Impacts on input use per sown manzana						
	(1)	(2)	(3)	(4)	(5)	(6)
	<u>ATE,</u>	<u>ATT,</u>	<u>ATE,</u>	<u>ATT,</u>	<u>ATE,</u>	<u>ATT,</u>
	<u>Fertilizer</u>	<u>Fertilizer</u>	<u>Chemical</u>	<u>Chemical</u>	<u>Organic</u>	<u>Organic</u>
	<u>Value/ha</u>	<u>Value/ha</u>	<u>Value/ha</u>	<u>Value/ha</u>	<u>matter kg/ha</u>	<u>matter kg/ha</u>
d (Participants in RBD=1)	9.591 [11.885]	17.638 [12.914]	99.967 [107.632]	116.496 [114.241]	1,522.04 [1,243.205]	581.721 [458.410]
Observations	243	243	243	243	243	243
R-squared (based on full sample)	0.1762	0.1681	0.1027	0.1232	0.2406	0.1338
Bootstrapped standard errors in brackets, 999 replications						

The results in Table 12 do not send a strong signal that the RBD program caused differences in how treated and untreated households used inputs in rice production. Columns (1) and (2) measure impacts on expenditures in Nicaraguan Córdobas on chemical fertilizers, Columns (3) and (4) represent impacts on expenditures on agricultural chemicals, and Columns (5) and (6) measure impacts on quantity of organic matter applied. None of the estimated impacts are significant, with the closest being the ATT for fertilizer with a p-value of 0.172, but even in that case the magnitude of the effect is very small (just over one U.S. dollar). It does appear that participant households applied inputs more intensively overall, as the sign of each coefficient is positive. But there is no detectable pattern that might indicate an impact on input composition that would drive yields downward. Beyond the direct effects of the drought, the exact mechanism by which the RBD program is failing to affect households or impacting them negatively remains unclear.

8. Conclusion

This paper has examined the Rural Business Development program for rice and plantain farmers administered by the Millennium Challenge Account in Nicaragua, a program designed to help small farmers become commercially viable producers. The descriptive component of the paper focused on both rice and plantain, and summary statistics based on survey data collected from participant and non-participant households indicate that there is far more heterogeneity between the two treatment groups in the case of plantain than with rice, at least with respect to wealth. Plantain farmers enrolled in the RBD program reported substantially higher asset values, both with respect to agricultural wealth (land, animals, and agricultural implements) and non-agricultural wealth (housing, consumer durables). In contrast, the characteristics of rice farmers do not appear to vary much with by RBD program participation status.

Following this preliminary examination of the data, I turned my attention to estimating impacts of the first year of the RBD program for rice farmers. A conservative conclusion based on these results would be that the program had no positive effects, whereas a conclusion that puts more faith in the assumptions of the estimated models would state that year one of the program caused large and negative impacts. It may be the case that the RBD program is focused on raising productivity, but when production conditions are sufficiently poor, no amount of assistance or access to high quality inputs is going to improve outcomes; in fact, these interventions may distract from usual risk management techniques and be detrimental as a result. But the data do not offer any insight into whether this is the truth. A fuller picture of how the RBD program affects rice production and household welfare will be painted once the second round of data are available, presuming climatic conditions have returned to normal.

References

- Angrist, J., Pischke, S. 2009. *Mostly Harmless Econometrics*. Princeton University Press, New Jersey.
- Brookhart, M., Schneeweiss, S., Rothman, K., Glynn, R., Avorn, J., Sturmer, T. 2006. "Variable selection for propensity score models." *American Journal of Epidemiology*, 163(12), 1149-1156.
- Caliendo, M., Kopeinig, S. 2005. "Some practical guidance for the implementation of propensity score matching." IZA Discussion Paper Series No. 1588.
- Greene, W. 2003. *Econometric Analysis*. Pearson Education, New Jersey.
- Hardle, W. 2007. *Applied Multivariate Statistical Analysis*. Springer Publishing, New York.
- Hirano, K., Imbens, G. 2002. "Estimation of causal effects using propensity score weighting: an application to data on right heart catheterization." *Health Services and Outcomes Research Methodology*, 2(1), 259-278.
- Imbens, G. 2004. "Nonparametric estimation of average treatment effects under exogeneity: a review." *Review of Economics and Statistics*, 86(1), 4-29.
- Robins, J., Rotnitzky, A., Zhao, L. 1995. "Analysis of semiparametric regression models for repeated outcomes in the presence of missing data." *Journal of the American Statistical Association*, 90(429), 106-121.
- Rosenbaum, P., Rubin, D. 1983. "The central role of the propensity score in observational studies for causal effects." *Biometrika*, 70(1), 41-55.
- Wooldridge, J. 2007. "Inverse probability weighted estimation for general missing data problems." *Journal of Econometrics*, 141(2), 1281-1301.

Appendix 1. Technical explanation of estimation methods

A1.1 Combining linear regression and inverse propensity score weighting

The following section is a technical description of the estimation methods used in this report, which may continue to serve as the basis of the evaluation strategy once the second round of data is collected. The evaluation strategy for both the plantain and rice components of the RBD program will rely on propensity score methods. For both the plantain program and the rice program once the second round of data collection is complete, these will be combined with panel data methods. There are a variety of methods that employ the propensity score as a means of generating consistent estimates of the ATE. Here, we will rely on inverse propensity score weighting (IPW); see Wooldridge (2007) for theory behind IPW in the context of treatment effect estimation, and Hirano and Imbens (2002) for an application. By constructing weights based on $p(x)$, we can correct imbalances in the distribution of X in the treated and control groups. Suppose we would like to estimate the ATE, and that as in the main text, the outcome of interest is denoted by y and the treatment indicator by the dummy variable d . We have:

$$E\left[\frac{dy}{p(x)}\right] = E\left[E\left[\frac{dy(1)}{p(x)} \mid X = x\right]\right] \quad (12)$$

This holds due to the law of iterated expectations and because $dy=dy(1)$. Since potential outcomes $y(1)$ and $y(0)$ are assumed to be independent of the treatment indicator d conditional on X , $E(dy(1) \mid X = x) = E(d \mid X = x)E(y(1) \mid X = x)$. Combining this with (12) yields:

$$E\left[E\left[\frac{dy(1)}{p(x)} \mid X = x\right]\right] = E\left[\frac{p(x)E(y(1) \mid X = x)}{p(x)}\right] = E[E(y(1) \mid X = x)] = E[y(1)] \quad (13)$$

which is the average value of y when everyone participates in the RBD program. Similarly, for the untreated outcome, we have:

$$E\left[\frac{(1-d)y}{1-p(x)}\right] = E\left[E\left[\frac{(1-d)y(0)}{1-p(x)} \mid X = x\right]\right] =$$

$$E\left[\frac{1-p(x)E(y(0) \mid X = x)}{1-p(x)}\right] = E[y(0)] \quad (14)$$

which is the average value of y when no one receives treatment. The difference between these two parameters is the ATE of the RBD program on y . In other words, we can recover the ATE of the RBD program by weighting the observed outcomes for treated households by the inverse of the propensity score and the untreated households by the inverse of $(1-p(x))$, calculating the respective group averages of the weighted outcomes, and taking the difference between these two averages. These weights can be expressed as:

$$\frac{1}{p(x)} \quad (15)$$

for treated households and

$$\frac{1}{1-p(x)} \quad (16)$$

for untreated households, or written as a single parameter:

$$\frac{d}{p(x)} + \frac{1-d}{1-p(x)} \quad (17)$$

Applying these weights to the treated and untreated outcomes in the population, we have the ATE:

$$E\left[\frac{dy}{p(x)}\right] - E\left[\frac{(1-d)y}{1-p(x)}\right] = E[y(1)] - E[y(0)] \quad (18)$$

This would be estimated using data by replace the true propensity score with an estimate, and taking sample averages instead of expectations. Suppose we estimate a model for the propensity

score and use it to generate predicted probabilities of participation. This would allow (18) to be estimated as:

$$\sum_{i=1}^N \frac{y_i}{N} \frac{d_i}{\hat{p}(x)} - \sum_{i=1}^N \frac{y_i}{N} \frac{1-d_i}{1-\hat{p}(x)} \quad (19)$$

Note that in order for (19) to be a consistent estimate of (18), we must correctly specify the model of the propensity score, i.e., $\hat{p}(x) = p(x)$. If it is correctly specified, then (19) converges to the ATE as the sample size N grows.

Instead of measuring the average impact of the RBD program on all households that satisfy program criteria, we might instead choose to estimate the ATT. The ATT is the difference between the average outcome in the treated group, and the average outcome that would have occurred in the treated group had none of them participated in the RBD program. This first component can be estimated by taking expectation of the outcome in the treated group, i.e.:

$$E[y(1) | d = 1] = E[dy] \quad (20)$$

where conditioning on $d=1$ indicates that the expectation is being taken over the subgroup of households that are members of the treated group, and each treated household is given a weight of 1; i.e., the weight given in (16) is multiplied by $p(x)$. For the second component of the ATT, we have:

$$\begin{aligned} E[y(0) | d = 1] &= E[E[y(0) | d = 1, X = x] P(d = 1 | X = x)] = \\ &E[E[y(0) | X = x] P(d = 1 | X = x)] = E[E[y(0) | X = x] p(x)] \end{aligned} \quad (21)$$

The second equality follows from the independence of $y(0)$ and d conditional on X . Note that the last term in (21) is equal to the first term in the second line of (14), multiplied by the propensity score $p(x)$. In other words, to recover the ATT, we adjust the weights given in (15) and (16) by multiplying through by $p(x)$, re-weight the observed outcomes, and then calculate the difference in the average outcomes for the treated and untreated groups.

We might not be very confident in our ability to correctly model the propensity score. As an alternative, we could use a regression-based approach that attempts to directly model the process determining the outcome variable y . For example, suppose that the true model for the average outcome with treatment conditional on $X=x$ is:

$$E[y(1) | X = x] = \alpha_0 + \alpha_1 + (x - \mu)' \beta_{treated} \quad (22)$$

And that the corresponding model for the untreated outcome conditional on $X=x$ is:

$$E[y(0) | X = x] = \alpha_0 + (x - \mu)' \beta_{untreated} \quad (23)$$

The parameter vectors $\beta_{treated}$ and $\beta_{untreated}$ capture the direct effect of the observables X on the outcome. The direct effects of household characteristics on the outcome might vary depending on whether or not a household participates in the program, and thus $\beta_{treated}$ and $\beta_{untreated}$ may differ. Subtracting the vector of means μ from observable characteristics X in the model has no effect on the values of $\beta_{treated}$ and $\beta_{untreated}$, but makes it easier to recover the ATE.

Taking the expected values of (22) and (23) over the distribution of X gives us:

$$E[y(1)] = \alpha_0 + \alpha_1 \quad (24)$$

and

$$E[y(0)] = \alpha_0 \quad (25)$$

These are the respective intercept terms of the equations given in (22) and (23). The difference between these intercepts yields the ATE as captured by α_1 . This model can be estimated by ordinary least squares (OLS), substituting sample averages for μ . An equivalent approach is estimating a single regression of y on X , d , and interactions between d and the de-meaned X variables, as was done in the paper.

It turns out that by combining IPW with the regression based approach we can take advantage of the “double robustness” property exhibited by IPW when combined with certain models, as originally shown by Robins et al. (1995) and demonstrated in the context of estimating the ATE by Wooldridge (2007). This requires estimating the model given in equations (22) and (23) by weighted least squares, where the weights depend on treatment status and are given by (15) and (16). Suppose the model specified above for the conditional expectations $E[y(1) | X = x]$ and $E[y(0) | X = x]$ is in fact correct. Estimating the parameters of the model by minimizing the weighted sum of squared residuals will yield consistent estimates of the parameters $\alpha_0, \alpha_1, \beta_{treated}$, and $\beta_{untreated}$, and thus yield a consistent estimate of the ATE. This holds regardless of whether or not the model for the propensity score is correctly specified; weighted least squares is always consistent when the regression model is correct (see Greene, 2003, page 226).

Now, suppose we have specified incorrect models for $E[y(1) | X = x]$ and $E[y(0) | X = x]$, but our model for the propensity score is correct, i.e., $\hat{p}(x) = p(x)$. The weighted least squares formula for the estimate of the intercept of (22) is:

$$\hat{\alpha}_0 + \hat{\alpha}_1 = \sum \frac{y}{N} \frac{d}{\hat{p}(x)} - \hat{\beta}_{treated} \sum \frac{(x' - \bar{X})}{N} \frac{d}{\hat{p}(x)} \quad (26)$$

And the formula for the estimated intercept of (23) is:

$$\hat{\alpha}_0 = \sum \frac{y}{N} \frac{1-d}{1-\hat{p}(x)} - \hat{\beta}_{untreated} \sum \frac{(x' - \bar{X})}{N} \frac{1-d}{1-\hat{p}(x)} \quad (27)$$

where the summations are taken over the sample from 1 to N . The correctly specified propensity score will insure that weighted averages of x in the treated and untreated subsamples will be

equal to the sample average \bar{X} ; i.e., the propensity score balances the distribution of the covariates in the treated and untreated observations. As a result the second summation terms in (26) and (27) drop out. Taking the difference of (26) and (27) therefore yields:

$$\sum_{i=1}^N \frac{y_i}{N} \frac{d_i}{\hat{p}(x)} - \sum_{i=1}^N \frac{y_i}{N} \frac{1-d_i}{1-\hat{p}(x)} \quad (28)$$

This is equivalent to (19), which we already proved to be a consistent estimate of the ATE when the propensity score is correctly specified.

Appendix 2. Description of data files

Table A2.1 below summarizes the contents of the data files, along with the relevant pages of the survey. All data files were created using Stata version 10.1. All files contain producer names, and the file “seccion_0_limpiar” includes GPS coordinates and altitude for land parcels.

Table A2.1. Summary of data files		
<u>File name</u>	<u>Survey pages</u>	<u>Description of contents</u>
seccion_0_limpiar	1-3	Identification of producer, RBD program eligibility criteria, geographic location.
seccion_i_limpiar	4	Housing characteristics
seccionii_ac_limpiar	5-6	Household roster and educational expenditures, 2009
seccion_ii_b_limpiar	5	Additional household members in 2008
seccion_iii_a_limpiar	7	Agricultural implements
seccion_iii_b_limpiar	7	Agricultural implements sold or lost in past year
seccion_iii_cd_limpiar	7	Animals, 2008 and 2009
seccion_iii_ef_limpiar	8	Consumer durables, owned at time of interview and sold or last in past year
seccion_iv_1_limpiar	10-12	Consumption in last 15 days
seccion_iv_2ae_limpiar	13-14	Consumption of transportation and health services in the past month
seccion_iv_2b_limpiar	14	Consumption of non-food items, past month
seccion_iv_2c_limpiar		Consumption of non-food items, past six months
seccion_iv_2d_limpiar		Consumption of non-food items, past year
seccion_v_limpiar	15	Transfers and remittances
seccion_vi_1_limpiar	16	Characteristics of owned land
seccion_vi_1d_limpiar	17	Sales and losses of land in past year
seccion_vi_1e_limpiar	17	Security of land tenancy
seccion_vi_2a_limpiar	18	Agricultural installations on owned land
seccion_vi_2b_limpiar	18	Improvements to owned land
seccion_vii_1a_limpiar	19	Experience with rice and production of rice in 2009
seccion_vii_1b_limpiar	20	Sales of rice planted in 2009
seccion_vii_1c_limpiar	21	Processing costs for rice, 2009
seccion_vii_1d_limpiar	21	Use of machinery, animal traction, and manual labor in land preparation and crop maintenance for rice, 2009
seccion_vii_1e_limpiar	22	Use of chemical and non-chemical inputs for rice, 2009
seccion_vii_1f_limpiar	22	Participation in the RBD program for rice, and other technical assistance
seccion_vii_2a_limpiar	23	Experience with plantain and production of rice in 2009
seccion_vii_2b_limpiar	23	Sales of plantain in 2009
seccion_vii_2c_limpiar	24	Use of machinery, animal traction, and manual labor in land preparation for plantain, 2009
seccion_vii_2d_limpiar	24	Use of machinery, animal traction, and manual labor in crop maintenance for plantain, 2009
seccion_vii_2e_limpiar	25	Use of chemical and non-chemical inputs for plantain, 2009
seccion_vii_2f_limpiar	25	Participation in the RBD program for plantain, and other technical assistance
seccion_vii_3a_limpiar	26	Area in other annual crops, 2009
seccion_vii_3b_limpiar	26	Area in perennials, 2008 and 2009
seccion_vii_3c_limpiar	27	Area in other annual crops, 2008
seccion_vii_3d_limpiar	27	Income from other businesses
seccion_vii_4a_limpiar	28	Formal credit loans, 2009
seccion_vii_4b_limpiar	28	Informal credit loans, 2009
seccion_vii_4c_limpiar	29	Formal credit loans, 2008
seccion_vii_4d_limpiar	29	Informal credit loans, 2008