NBER WORKING PAPER SERIES

HETEROGENEOUS IMPACT DYNAMICS OF A RURAL BUSINESS DEVELOPMENT PROGRAM IN NICARAGUA

Michael R. Carter
Emilia Tjernström
Patricia Toledo

Working Paper 22628
http://www.nber.org/papers/w22628

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2016

We thank Anne Rothbaum, Lola Hermosillo, Jack Molyneaux, Juan Sebastian Chamorro, Carmen Salgado, Claudia Panagua, Sonia Agurto, the staff at FIDEG, and Conner Mullally. We gratefully acknowledge funding from the Millennium Challenge Corporation, as well as financial support from the US Agency for International Development Cooperative Agreement No. EDH-A-00-06-0003-00 through the BASIS Assets and Market Access CRSP. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Heterogeneous Impact Dynamics of a Rural Business Development Program in Nicaragua
Michael R. Carter, Emilia Tjernström, and Patricia Toledo
NBER Working Paper No. 22628
September 2016
JEL No. I32,O12,O13,Q12,Q13

ABSTRACT

We study the impacts of a rural development program designed to boost the income of the small-farm sector in Nicaragua. Exploiting the random assignment of treatment, we find statistically and economically significant impacts on farm incomes and investment in farm capital. Using continuous treatment estimation techniques, we examine the evolution of program impacts over time and find that incomes in the activities targeted by the program as well as farm capital rise significantly over time, even after the expiration of the program. Because of the temporal pattern of impacts, shorter-term binary treatment estimators do not fully capture the impacts of the program. Additionally, panel quantile methods reveal striking heterogeneity of program impacts on both income and investment. We show that this heterogeneity is not random, and that there are some low-performing households that simply do not benefit from this program that tried to engage them as agricultural entrepreneurs. While the benefit-cost ratio of the program is on average highly positive, these findings on impact heterogeneity signal limitations of business development programs as a way to eliminate rural poverty.

Michael R. Carter
Department of Agricultural and Resource Economics
University of California, Davis
One Shields Avenue
Davis, CA 95616
and NBER
mrcarter@ucdavis.edu

Emilia Tjernström
Robert M. La Follette School of Public Affairs
University of Wisconsin–Madison
1225 Observatory Drive
Madison, WI 53706–1211
emilia.tjernstrom@gmail.com

Patricia Toledo
Department of Economics
Bentley Annex 360
Athens, Ohio 45701
toledot@ohio.edu

Heterogeneous Impact Dynamics of a Rural Business Development Program in Nicaragua

With severe poverty concentrated in rural areas of the developing world, there have been numerous efforts to engage the rural poor as entrepreneurs. The hope is that with the right information, investment and market connectivity, the poor can boost their incomes, invest in their children and work their way out of poverty. However, in contrast to cash transfer programs, which address poverty by "just giving money to the poor" (Hanlon, Barrietnos and Hulme (2010)), these business development programs that consider the poor incipient entrepreneurs exhibit several characteristics that shape their effectiveness and challenge the evaluation of their impacts:

1. Dynamics: By providing new information, incentives and connections, we might expect entrepreneurially-focused programs to induce beneficiaries to learn and to co-invest in their new opportunities, therefore making it likely that impacts will evolve over time.[1]

2. Participation: While most people can and do accept a cash transfer if one is offered, entrepreneurial programs require specialization, investment and risk-taking and are thus unlikely to appeal to all poor households. This limits their effectiveness as an anti-poverty strategy.

3. Impact Heterogeneity: Any entrepreneurial activity will generate both winners and losers, based on luck and/or complementary inputs that differ across households (*e.g.*, talents and skills), again limiting the average effectiveness of programs that address the poor as potential entrepreneurs.

While studies of other programs that address the poor as entrepreneurs have noted that partial participation blunts program impacts (e.g., see Banerjee et al. (2011)), this paper uses data from a 5-year study of a Nicaraguan program that was randomly rolled out over time to explore all three of these dimensions of addressing the rural poor as incipient agricultural entrepreneurs, ,

Nicaragua, one of the poorest countries in the western hemisphere, is no exception to the pattern in which poverty is most severe in rural areas. Beginning in 2007, the government of Nicaragua launched a rural business development program (RBD) in cooperation with the Millennium Challenge Corporation (MCC), the United States government foreign aid agency. The RBD was designed to address a set of constraints that were believed to restrict the productivity and incomes of resource-scarce rural households. Specifically, the RBD offered marketing interventions, temporary input subsidies and/or co-investment incentives, and

[1] As King and Behrman (2009) point out, programs with significant learning and adoption components are unlikely to attain steady-state effectiveness soon after an intervention begins. In this study, we therefore pay particular attention to how the observed impacts evolve over time.

extension services. Contact with farmers generally lasted 24 months, after which farmers were expected to continue on with their own knowledge and resources.

While none of these interventions are novel, prior non-experimental efforts to evaluate similar programs' effectiveness have confronted identification problems because of endogenous program placement and participation (see e.g. Evenson (2001) and Anderson and Feder (2003)). Several recent studies employ experimental designs to solve these identification problems: Bardhan and Mookherjee (2011), Carter, Laajaj and Yang (2013) and Carter, Laajaj and Yang (2014) find positive impacts of subsidized agricultural inputs to farmers in West Bengal and Mozambique, respectively, and Cole and Fernando (2012) find that farmers respond to mobile-phone based agricultural information delivery in Gujarat. Ashraf, Giné and Karlan (2009) use an experimental design to estimate the impact of extension services and find positive early impacts on incomes.[2] However, unlike Carter, Laajaj and Yang (2014) who find that positive impacts evolve but persist over time, the impacts in Ashraf, Giné and Karlan (2009) completely dissipate over time, reinforcing the importance of paying attention to impact dynamics.

To evaluate the impacts of the Nicaraguan RBD program, we worked with program implementers to select a random subset of program-eligible households for inclusion in the study. These study households were in turn randomly split into early and late treatment groups, as the treatment could not be rolled out to all households at once due to capacity constraints on the implementer side. Early-treatment households were offered the program in 2007, shortly after a baseline survey was conducted. The late-treatment households were offered the program some 20 months later, after the second (mid-line) survey. A third (end-line) survey took place another two years on, in 2011. The result is a 3-round panel data set, in which final exposure to the program randomly varies across households from as much as 4 years to as little 18 months.[3] We exploit the fact that the late-treament households made their program participation decisions after the mid-line survey, which allows us to realize statistical efficiency gains by focusing the analysis only on those who participate in the program (a double-complier sample). Further, while the baseline and mid-line data have a conventional binary-treatment/control structure, the full 3 rounds of the panel data allow us to use fixed-effect continuous treatment estimators to trace out program impacts over time.

Using this design, we explore RBD's impacts on three key outcome variables: income in targeted agricultural activities, productive investment, and per-capita household consumption expenditures. We find significant average impacts of the RBD on income and investment, but not on household consumption expenditures. Our estimates show that the impacts evolve over time and suggest that the standard binary

[2] See also Feder, Slade and Lau (1987) for an earlier study of extension service intensification using a quasi-experimental research design, which uncovers positive but diminishing effects of extension services.

[3] While most of the variation in treatment duration is between early and late groups, a small portion also results from variation within-groups. While we did not randomize within-group treatment order, there is no evidence (qualitative nor quantitative) that it was anything but random.

treatment estimates based on the mid-line data present an incomplete picture of long-term impacts. In particular, the average impacts of the RBD program on farm-level capital investments continue to grow after the mid-line survey, suggesting that longer time-frames may be necessary to appropriately evaluate these types of programs. The failure of consumption expenditures to respond to the RBD program appears to reflect households' decisions to reinvest income increases rather than consume them.

Looking beyond average impacts, we employ the panel quantile regression techniques developed by Abrevaya and Dahl (2008) to determine the extent to which estimated average impacts represent the range of impacts experienced by program participants. The analysis reveals quite striking heterogeneity in impacts. Beneficiaries in the 75^{th} conditional quantile of incomes enjoy much larger impacts than those in the lower quantiles. Indeed, program impacts are insignificant for those in the lowest quantiles. and we estimate that those in the top conditional quantiles enjoy income and investment impacts that are roughly double the median impacts. Not surprisingly, the average impact paths appear steeper than those estimated by a median regression, as they are driven up by the OLS regression's sensitivity to extreme values.

While there are multiple explanations why impacts may be nil in the lower quantiles, we present evidence to suggest that there is relatively little movement of households across quantiles over time. That is, there appear to be "lower quantile" type households who benefit little from the RBD program, and high types who benefit substantially. This finding, along with a 60% program participation rate suggests that the RBD is an effective tool for raising incomes for some: it places a strong minority of households on an upward economic trajectory. However, it also appears to be an ineffective tool for many others.[4]

The remainder of this paper is organized as follows. Section 1 introduces the RBD, describes the data, and presents basic descriptive statistics and balance tests between the early and late treatment groups. Section 2 presents the basic econometric strategy, with section 3 showing the average impact estimates for income, investment and consumption. Section 4 looks beyond average impacts and estimates the extent of impact heterogeneity and its meaning. Section 5 concludes.

1 Background

Agriculture has played an important role throughout Nicaragua's history, but it is widely held that multiple constraints have conspired to prevent agriculture from reaching its productive potential—examples include a lack of basic infrastructure, low education levels, and low access to credit and technology. The Western Region of Nicaragua, which includes the departments of León and Chinandega, was identified by a National

[4]By way of comparison, Banerjee et al. (2011) find that approximately one-third of intended beneficiaries declined participation in a business development program that offered a free asset transfer. These authors do not, however, break down the distribution of beneifits across household types. None of these observations mean that programs like the RBD are bad policy, simply that by themselves they are inadequate to raise the living standards of all targeted households.

Development Plan (NDP) as having particularly high potential for agricultural growth. While high-potential, the area is also quite poor: the World Bank (2008) determined that more than 50 percent of households in the Western Region live in poverty.

In July 2005, the Millennium Challenge Corporation (MCC) signed a five-year, $175-million compact with the Government of Nicaragua to develop a set of projects in the Western Region, with the objective of relaxing some of the aforementioned constraints. The compact had three components: a transportation project, a property regularization project, and the one we focus on here: a rural business development (RBD) project.[5] This latter component aimed to raise incomes for farms and rural businesses by helping farmers develop and implement a business plan built around a high-potential activity, as explained in more detail below.

1.1 Program Description

The Nicaraguan implementing agency (the Millennium Challenge Account, or MCA) identified the productive activities most suitable for inclusion in the program: beans, cassava, livestock, sesame, and vegetables. In order to be eligible, farmers had to own a small- or medium-sized farm, have some experience with one of these crops, be willing to develop a business plan together with extension agents, and contribute 70% of the cost of investments identified in the business plan. In addition, activity-specific eligibility criteria were developed and applied (the precise rules are shown in Appendix A).[6] Once farmers enrolled in the program and had their business plan approved, the RBD program worked with them for 24 months. While the exact benefits varied across the productive activities, all farmers received technical and financial training and supplies based on their individual business plan. One-time input subsidies for improved seeds and fertilizers were also provided in some cases.

Farmers were grouped into small geographical clusters of approximately 25 farmers, with a lead farmer identified for each. The randomization exploited the fact that capacity constraints meant that not all eligible farmers could be brought into the project immediately. The research team worked with the RBD implementers to identify all the geographical clusters that would eventually be offered RBD services. The evaluation team then selected a subset of these clusters for random assignment to either *early* or *late* treatment status.

Once the random assignment of early and late clusters had been finalized, 1,600 households were sam-

[5] The MCC terminated a portion of the compact in June of 2009, reducing compact funding from $175 million to $113.5 million. While this action cut off the property regularization part of the program, the RBD Program was not affected by this partial project termination.

[6] The impact of these eligibility criteria on the characteristics of the eligible population is described in Toledo and Carter (2010) who show that the RBD beneficiaries are found in the middle deciles of the rural income distribution of the areas where the program was implemented (the Departments of León and Chinandega).

pled from the roster of all eligible producers in these clusters, split equally between early and late areas. Approximately 12 farmers were randomly selected for the study from each cluster. These 1,600 households completed a baseline survey in late 2007, just as the RBD program was rolling out in the early treatment clusters. The mid-line survey was conducted approximately 18 months later, right before the late treatment group was offered the program. As illustrated in Figure 1, the randomization and the timing of the surveys meant that the late treatment group function as a temporary control group at the time of the mid-line survey. Both early and late treatment clusters were then surveyed again near the end of the program in 2011. This roll-out strategy also effectively randomized the duration of time in the program, a feature that will prove improtant in the continuous treatment estimates presented below.

Figure 1: Timeline of Received Treatment and Timing of Surveys

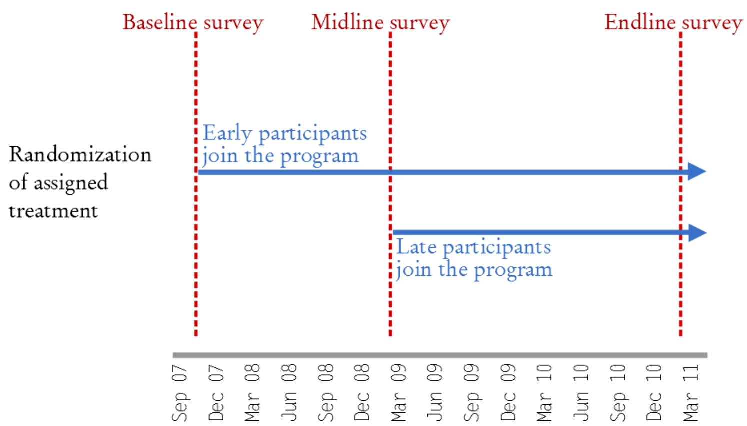

Out of the eligible households in the early (late) treatment clusters, 64% (57%) chose to participate in the RBD project. While we could not foresee who would reveal themselves to be compliers at the baseline, the late-treatment households made their participation decision around the time of the mid-line survey. We can therefore estimate the impacts both on *eligible* households and on participating or *complier* households, and Section 3 shows that the results are very similar, suggesting that the decision to take up the program was consistent in the early and late treatment groups.

1.2 Data

As the previous section explained, we have a three-wave panel of 1,600 households with less than 2 percent attrition by the time of the third wave. Table 1 shows summary statistics and two separate balance checks. Because our control (late treatment) group was eventually offered the program, we check for statistical balance in two ways: the first set of columns in the table compare means for households randomized into the early and the late treatment groups, while the second set compare the self-selected complier households in the early and late groups.

Table 1: Summary Statistics and Balance Checks

	Full sample			**Compliers**		
	Early	Late	Diff.	Early	Late	Diff.
Household characteristics						
	(1)	(2)	(3)	(4)	(5)	(6)
Program farmer: age (#)	52.48	50.21	2.27***	52.63	49.41	3.22***
	(12.88)	(13.35)		(12.63)	(13.13)	
Program farmer: education (#)	3.99	4.22	-0.24	3.95	4.44	-0.49*
	(4.16)	(4.31)		(4.09)	(4.33)	
Program farmer: gender (=1 for female)	0.13	0.14	-0.01	0.13	0.14	-0.0172
	(0.34)	(0.35)		(0.33)	(0.35)	
Household members (#)	5.46	5.33	0.13	5.51	5.47	0.036
	(2)	(2)		(2)	(2)	
Per capita expenditures ($)	3998	3815	183	3868	3806	62.26
	(5157)	(4102)		(4074)	(4276)	
Farm characteristics						
Value of mobile capital ($)	3590	4290	-700.4	3620	4502	-882.1*
	(7350)	(10965)		(6836)	(9415)	
Value of fixed capital (installations) ($)	3148	3219	-70.82	2966	3504	-538.2
	(6482)	(6901)		(4633)	(8042)	
Years of experience in target crop	21.28	21.19	0.09	21.71	21.23	0.48
	(13.13)	(13.03)		(13.33)	(12.45)	
Landholdings: owned (manzanas)	41.85	34.57	7.28**	42.42	38.41	4.01
	(81.12)	(53.28)		(74.12)	(61.34)	
Landholdings: amt. planted in target crop	4.93	5.21	-0.28	5.61	5.12	0.48
	(7.50)	(10.46)		(8.55)	(6.61)	
Landholdings: amt. planted in maize	2.96	3.10	-0.14	3.00	3.22	-0.22
	(2.03)	(2.74)		(2.01)	(3.11)	
Share of seasons used improved seeds (=1 for improved)	0.13	0.12	0.01	0.15	0.11	0.04*
	(0.33)	(0.32)		(0.36)	(0.31)	

Standard errors in parentheses
The asterisks in the third and fourth columns denote the statistical significance of
t-tests on the equality of means between early and late groups with asterisks indicating significance:
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Most of the variable averages suggest that the groups represent the same population, (*i.e.,* they are not statistically significantly different from each other) and that the randomization was successful. Comparing all eligible farmers, farmers in the early group are slightly older than those in the late group, and they own slightly more land. The difference in age holds also in the sample of compliers sample, but the actual age difference is small enough that we doubt that it would be economically significant, especially given that the years of experience in the target crop are very similar across the two groups.

The difference in landholdings might be worrisome, as it is a key productive asset, but if we look instead at the amount of land planted in the target crop, the groups plant very similar amounts of land in the target activity. Further, once we examine the complier sample, the differences in landholdings between the two groups are no longer statistically significant. However, a t-test indicates that the value of mobile capital is statistically different between the two groups. Taken together, some of the statistically significant differences between treatment and control households imply that early treatment households are better off (landholdings), while others imply that they are worse off (mobile capital). Combined with the fact that the statistical significance switches on and off in the full and the complier sample, we believe that the differences are of the kind that one would expect by accident even after a successful randomization. Nonetheless, our preferred specifications adopt a difference-in-difference strategy rather than simple differences to ensure that any baseline differences do not drive our results.

Figure 2 shows a histogram detailing the distribution of months in the RBD program for the sample of compliers across all three survey rounds. The figure excludes observations with 0 months of treatment since this group (comprised of the early and late treated households at baseline, plus the late treated households at the mid-line), dwarfs the other categories. Despite some bunching, the data show reasonable dispersion: the data contain households observed with as little as 1 month in the program up to as much as 50 months in the program. The largest clusters of observations are around 6, 18 and 40 months of program exposure. Late treatment households comprise the first group, both early and late treatment households are found in the middle exposure group (the former at mid-line, the latter at end-line), while the latter group is comprised exclusively by early treatment households. It is this variation in length of time in the program that will be exploited in the continuous treatment estimators explained in the next section.

2 Econometric Methodology

Our three outcome variables of interest – farm income, investment, and household consumption – capture both direct and indirect channels of impact. The small-farm intervention was designed to enhance the access of small farmers to improved technologies and to markets, so we begin by examining program impacts on income in the target crops. We define income as the total value of production in the target crop, calculated using the prices that the household obtained for the part of their harvest that was sold.[7] While this measure is likely an upper bound on the total income impacts of the program, we can examine the effects of program participation on maize production, to examine whether the target crops are crowding out production of the basic staple crop in the target area. Section 2.2 shows that there is indeed some substitution of land away from maize and into the target crop. It is therefore worth keeping in mind that any income increases in the target activity are upper bounds on the income impacts of program participation.

[7]Note that the RBD was intended to allow farmers to receive better prices for their produce, hence it is important that we value output based on prices actually received. When a farmer did not sell any part of their crop, we valued output using the mean price in their geographical cluster by season and crop.

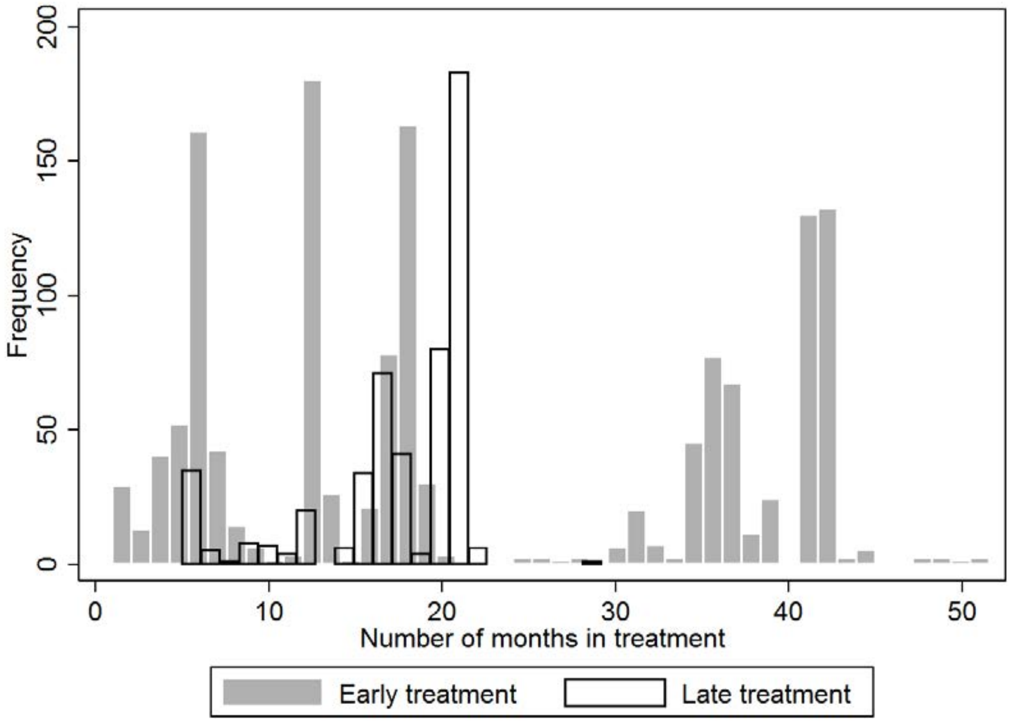

Figure 2: Distribution of the Duration of RBD Treatment (Dual Complier Sample) – Excluding Pre-treatment Observations

If the program results in increased agricultural income, households would then face a key choice: allocate income increases immediately to consumption, or reinvest the income into the farm operation, postponing increased consumption until a later date. We therefore examine investment and household consumption in turn.

We evaluate the impacts of the program using two main econometric approaches. First, we estimate the local average treatment effect (LATE) using Analysis of Covariance (ANCOVA) estimation on two different samples. The standard approach, given that the treatment was randomly assigned, would use treatment assignment to instrument for treatment status to compute the Treatment on the Treated (ToT) estimates. We compare this approach to a two-sided complier (2SC) estimator, which is similar to the standard ToT approach, but that allows us to gain power compared to standard approaches. Second, we employ a fixed effects, continuous treatment estimator to examine the evolution of impacts over time.

To motivate our focus on continuous treatment effects, note that the workhorse impact evaluation estimators assume that program participation is a binary state–either a household receives the treatment or it does not. While this approach deals well with treatment heterogeneity across treated units (hence the derivation of local average treatment effects), it is not equipped to deal with impacts that evolve over time.

Programs like the RBD that provide information, improve market access and enhance investment incentives might be expected to achieve their full impact over a medium-term time period of unknown duration. In the extreme case, they may even cause short-term decreases in key indicators as households switch livelihood strategies or even cut consumption to fund investments (Keswell and Carter (2014) find evidence of these short term dips in the case of land redistribution in South Africa).

Figure 3: Hypothetical Impact Patterns

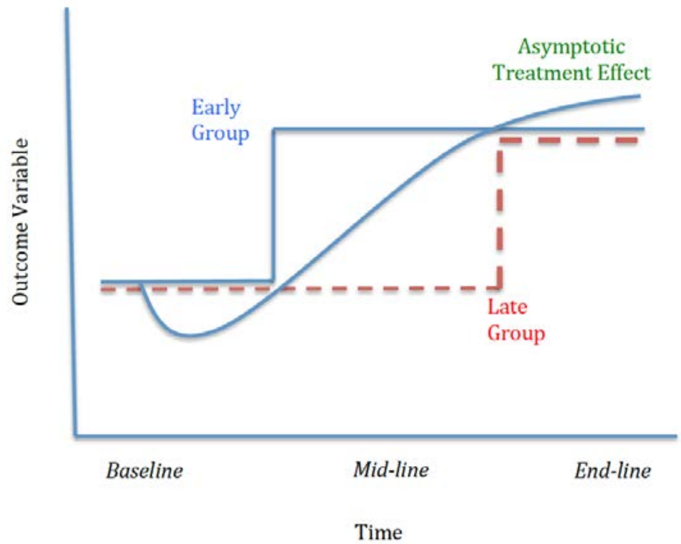

To better frame these issues, consider the hypothetical impact relationships for the RBD intervention illustrated in Figure 3. The solid, blue curve illustrates what we might expect to see for the early treatment group, while the dashed red curve illustrates the same for the late treatment group. The horizontal axis shows approximately the time at which the different survey rounds were undertaken. If the program had reached its full long-term impact on the early (late) treatment group by the time of the second (third) round survey, then conventional binary estimators would work well. In this case we would expect the data to trace out impact patterns similar to the step functions.

On the other hand, if the impact of the program evolves more slowly over time (for example, with an initial dip followed by a slow rise toward a long-run or asymptotic treatment effect), then our data would be generated by a a non-linear impact or duration response function in which impact depends on the duration of time in the program. Impacts measured at mid-line using a conventional binary treatment estimator (which works well when the data follow a pattern shown by the step functions in Figure 3) would reveal muted effects that would not accurately represent the long-run program impacts. The remainder of this section develops both binary impact estimators as well as a more general continuous treatment model designed to capture an unknown impact pathway.

2.1 Binary Treatment Model

In the binary analysis, we use ANCOVA estimation for the basic treatment estimates. McKenzie (2012) demonstrates that ANCOVA estimation can result in substantial improvements in power compared to the more common difference-in-difference specifications. The power gains are especially large when the data have low autocorrelation, as is the case for many outcomes in rural development settings like ours.

We begin by defining two indicator variables:

- B_i indicates treatment assignment for household i, equaling 1 for eligible farmers who were assigned to the early treatment group, and 0 for those assigned to the control or late treatment group.

- D_i indicates whether or not a farmer actually participated in the program when invited, so that $D_i = 1$ for treated invited farmers and $D_i = 0$ for non-compliers, who refused the program.

The local average treatment effect (LATE) can then be estimated by the coefficient δ in the instrumental variables ANCOVA regression:

$$y_{i,2} = \alpha + \delta \hat{D}_i + \theta y_{i,1} + \beta' X_{i,1} + \varepsilon_i \tag{1}$$

where where $y_{i,2}$ is the outcome variable in the second (post-treatment) period, $\hat{D}_i$ is D_i instrumented by B_i (the assignment to early treatment), $y_{i,1}$ is the baseline, pre-intervention value of the outcome variable for household i, and $X_{i,1}$ is a vector of baseline variables for which we want to control. Since the intervention was randomly assigned, the use of B_i as an instrument for D_i allows us to obtain consistent estimates of δ. We will present our results both with and without covariates, since the intervention was randomized.

Looking ahead to the continuous treatment model, where we only observe duration of time in treatment for the compliers (households with $D_i = 1$), we also employ a two-sided complier estimator, which instead of instrumenting for program take-up restricts the sample to the complier sample, i.e. farmers in both early and later groups who joined the RBD program. We are able to do this thanks to our third round of data, in which we observe the take-up decisions of the late treatment group, *i.e.* those farmers who serve as controls in the midline survey. In this case, the vast majority of program costs were spent on participating farmers, so the estimated impacts on this subpopulation are likely the most relevant to policymakers.

The estimating equation for the 2SC estimator is the same as in Eq. 1, except that instead of instrumenting for D_i using treatment assignment, we use the information from the third survey round to identify the compliers among the late treatment group.

The validity of this 2SC estimator relies on the idea that the decision to enroll in the early and late treatment groups was structurally the same, so that we are in fact comparing like-with-like in using this estimator. This assumption is in addition to the usual exclusion assumption, i.e. that farmers who do not enroll in the program experience no effect from the treatment or the randomization. As we show in Section 3, the estimated coefficients are very close whether we instrument or not, giving us some confidence that comparing compliers to compliers in the continuous analysis is sensible.

2.2 Continuous Treatment Model

As discussed in the beginning of this section, there are a number of possible reasons why the impact of the RBD program may have evolved over time. In addition to a possible initial dip in living standards when households first join the program and focus their resources on building up the targeted activity, there are at least three other reasons why the impact of the small-farm intervention may have changed over time. First, program beneficiaries may have experienced a learning effect, with their technical and entrepreneurial efficiency improving over time. Second, the asset program may have created a crowding-in effect if the program incentivized beneficiaries to further invest in their farms. As Keswell and Carter (2014) discuss, these second-round multiplier effects are what distinguish business development and asset transfer programs from cash transfer programs and other common anti-poverty policy instruments. Third, and less positively, if program impacts are short-lived (*e.g.*, if treated farmers drop the improved practices as soon as the 24-month period of intense RBD involvement with their groups end), then impacts may dissipate over time.

One goal of this study is to estimate the impact dynamics and duration response function, and thus recover both the long-run impacts of the intervention and their time path. Both are of particular relevance from a policy perspective. Indeed, it is the prospect that a skill-building program like the RBD program will facilitate and crowd-in additional asset building that makes them especially interesting as an anti-poverty program.

We begin by generalizing the binary response function to the continuous treatment case:[8]

$$E[y_{it}|d_{it}] = \alpha_i + \tau_2^d t_2 + \tau_3^d t_3 + f(d_{it}), \tag{2}$$

where d_{it} is the number of months since farm i was actively enrolled in the treament at survey time t, t_2 and t_3 are time dummies, and $f(d_{it})$ is a flexible function that can capture the sorts of non-linear impacts illustrated in Figure ((3)) above. These durations run from 0 to 50 months.[9]

[8] We could alternatively follow the generalization of propensity score matching to the continuous treatment case found in Hirano and Imbens (2004). The Hirano and Imbens estimator only exploits observations with strictly positive amounts of treatment. In our case, this would imply dropping the baseline data for all RBD participants as well as the mid-line data for the late treatment group. For development applications that employ this estimator, see Keswell and Carter (2011) and Aguero, Carter and Woolard (2010)

[9] In a few cases, RBD activities began a few months prior to the baseline survey. For these cases, we have considered households in these clusters as treated at baseline, but their values for d_{it} can exceed the number of months between the first

Based on the semi-parametric estimates of 2 reported in Tjernström, Carter and Toledo (2013), we choose a cubic parametric form to represent the duration impact function, $f(d_{it})$. The household-specific fixed effect term, α_i, controls for all observed and unobserved time-invariant characteristics, including farming skill, soil quality, farmer education, *etc.* Importantly, the fixed effect estimator controls for any systematic or spurious correlation between time invariant household characteristics and duration of treatment.

While there are several computationally equivalent ways to consistently estimate a fixed effect model like equation 2, in anticipation of later quantile regression analysis (where such models are less easily estimated), we will build on the correlated effects model of Mundlak (1978) and Chamberlain (1982,1984) and write the individual fixed effects as a linear projection onto the observables plus a disturbance:

$$\alpha_i = \lambda_0 + X_{i1}^{'}\lambda_1 + X_{i2}^{'}\lambda_2 + X_{i3}^{'}\lambda_3 + \upsilon_i,$$

where X_{it} denotes a vector of observables, which includes the time dummies and the duration variables. In our case, we have little reason to believe that the way in which the time-varying observables affect the individual effects differ between survey rounds, so we use the average of the time-varying covariates and write the fixed effect as

$$\alpha_i = \lambda_0 + \bar{X}_i^{'}\bar{\lambda} + \upsilon_i.$$

Substituting this expression into (2) gives:

$$y_{it}(d_{it}) = \tau_2^d t_2 + \tau_3^d t_3 + f(d_{it}) + \lambda_0 + \bar{X}_i^{'}\bar{\lambda} + [\upsilon_i + \varepsilon_{it}] \tag{3}$$

where ε_{it} is the error associated with the original regression function, equation 2. Replacing $f(d_{it})$ with the cubic functional form suggested by the semi-parametric analysis yields:

$$y_{it}(d_{it}) = \tau_2^d t_2 + \tau_3^d t_3 + \zeta_1 d_{it} + \zeta_2 d_{it}^2 + \zeta_3 d_{it}^3 + \lambda_0 + \bar{X}_i^{'}\bar{\lambda} + [\upsilon_i + \varepsilon_{it}]. \tag{4}$$

OLS estimation of (3) allows us to consistently recover the fixed effect estimators of the impact response function parameters of interest.

3 Average Impact Estimates

Using the binary and continuous treatment models developed above, this section presents estimated average RBD impacts for each of our three primary outcome variables: gross income in the targeted business

and third rounds of data collection.

activity, productive investment, and household living standards as measured using typical living-standards measurement survey consumption expenditure modules. Section 2.1 presents binary results using both the full sample and the 2SC estimator that restricts the sample to complier households. The 2SC complier estimates are strikingly similar to the IV estimates, suggesting that the compliers in the late treatment group are similar to those in the early treatment group and confirming that the program was carried out in a similar fashion for the two groups. Further, Table 2 shows a regression of the probability of program take-up decision on early/late treatment. Column (1) includes only the treatment assignment dummy, while column (2) adds in crop fixed effects and column (3) adds in other covariates. Treatment assignment is not statistically significant in any of the regressions, lending further support to the notion that the take-up decision similar in the two groups. Section 3.2 then presents results for the continuous treatment model, which uses only the complier sample. The similarity of LATE and 2SC complier results in this section suggests that this is a reasonable approach.

Table 2: Decision to Take Up Program, by Early and Late Treatment Group Status: Probit regression

	(1)	(2)	(3)
Early treatment (0/1)	-0.12	-0.13	-0.13
	(0.11)	(0.098)	(0.097)
Female household head			0.034
			(0.092)
Age of household head			-0.0032
			(0.0031)
Education of household head			0.0072
			(0.0087)
Household size			0.034**
			(0.014)
Landholdings (manzanas)			0.00087
			(0.00088)
Share of seasons used improved seeds (=1 for improved)			0.069
			(0.11)
Constant	0.51***	0.59***	0.51**
	(0.091)	(0.12)	(0.23)
Crop fixed effects?	NO	YES	YES
Pseudo-R^2	0.0018	0.013	0.019
N	1600	1600	1600

Clsuter-robust standard errors in parentheses
* $p<.1$, ** $p<.05$, *** $p<.01$

3.1 Binary impact estimates

Farm Income

We begin the evaluation of the RBD program by looking at its impact on annual farm income from the activities targeted by the RBD program.[10] As discussed above, observed income increases in RBD-targeted crops do not necessarily imply increased overall incomes, as productive inputs may have been reallocated from other activities, such as maize, to the target crops. Maize was not an RBD activity, but it is an important staple crop that most households produce. We do not have reliable maize prices, so we can only examine the amount of land devoted to maize versus the productive activity. This will provide a sense of whether intensification in the target crops comes due to improved technologies, or from a substitution effect thatc comes at the expense of reduced output and income from other crops.

Table 3: Impact of RBD Program on Target Activity Income: ANCOVA Estimates

	ITT			LATE		LATE (Complier sample)		
Treatment	897.7*	1003.0*	940.9	1360.5*	1502.7*	1345.0**	1495.6**	1437.9**
	(527.0)	(543.1)	(572.1)	(802.7)	(816.8)	(641.3)	(649.9)	(683.5)
Baseline farm income	0.84***	0.82***	0.82***	0.84***	0.82***	0.84***	0.90***	0.90***
	(0.072)	(0.080)	(0.080)	(0.070)	(0.078)	(0.072)	(0.079)	(0.079)
Education, household head		61.9	61.8		55.6		-62.8	-62.9
		(73.9)	(73.9)		(73.5)		(92.7)	(92.5)
Experience with target crop		-8.42	-8.26		-8.34		-13.9	-13.6
		(21.6)	(21.6)		(21.1)		(27.6)	(27.6)
Household size		140.0	139.7		124.1		101.0	101.1
		(109.6)	(109.7)		(110.7)		(153.8)	(154.1)
Land owned (manzanas)		13.9*	13.8*		13.6*		15.7**	15.6**
		(8.25)	(8.26)		(8.08)		(7.43)	(7.44)
Share of improved seeds		671.5	676.4		727.3		701.2	703.2
		(794.4)	(795.7)		(796.7)		(1169.6)	(1170.3)
Female household head		-376.9	-654.1		-423.0		-126.7	-370.0
		(469.8)	(686.0)		(452.2)		(563.6)	(721.1)
(Female head) x (treatment)			498.1					445.5
			(1018.9)					(1116.2)
Constant	295.1	-807.0	-771.2	257.0	-726.7	15.8	-410.9	-385.4
	(364.3)	(789.5)	(786.9)	(375.7)	(756.7)	(439.4)	(1173.0)	(1173.2)
N	1495	1433	1433	1495	1433	1028	993	993
$\overline{R}^2$	0.588	0.589	0.589	0.591	0.592	0.625	0.629	0.629

Cluster-robust standard errors in parentheses

* p<.1, ** p<.05, *** p<.01

All regressions control for crop fixed effects. Share of improved seeds measures the share of seasons at baseline that the beneficiary household used improved seeds

Table 3 shows the results from ANCOVA regressions on income at midline, with Intention-to-Treat (ITT)

[10] RBD targeted activities are beans, sesame, or cassava for farmers in those groups, and milk for livestock farmers. Income from these activities is the total value of production in the targeted activity, valued in 2005 $USPPP. We exclude vegetable farmers, as we were unable to obtain good price information on the many different vegetable crops.

estimates in columns (1) - (3), and LATE estimates in columns (4) - (5). These are the standard impact estimates under randomized treatment assignment, and make no assumptions about the uptake rate in the late treatment group, who act as control group in the midline. Further, columns (6) - (8) show the LATE results when we restrict the sample in both early and late treatment group to compliers only.

The ITT estimates show substantial average impacts of the program. The point estimates in our preferred regression in column (2) are roughly $1,000, and are significant at the 10-percent level. Column (3) reports a specification in which we interact the gender of the farmer with the treatment variable in the regression. As the descriptive statistics in Table 1 show, almost 90% of the individuals reported to be the main beneficiary of the RBD program are men. The estimated coefficient is quite large in absolute value (the impact of the program is almost $500 higher for women than for men), but it is imprecisely estimated–likely due to the fact that few women were actually enrolled in the RBD program. The results from the 2-sided compliance (2SC) estimator are very similar to the results in columns (4) and (5), which were obtained using standard instrumental variables regression in the full sample. The main difference is that the 2SC estimates are statistically significant at the 5-percent level; this increase in precision comes from not having to instrument for uptake.

Farm Capital

An important objective of beneficiaries' business plans was the accumulation of farm assets. With the objective of increasing farmers' productivity, the program provided some equipment or supported the construction of new productive installations once the business plan was approved. We follow the same strategy used in the previous section to examine the effect of the program on stocks of capital. The outcome variable used here is the sum of investments in mobile capital (tools and equipment, excluding livestock) and in fixed capital (buildings, installations, and fences located on the farmer's land).[11] The results are similar if disaggregated by type of capital. Note that in contrast to the income analysis, these measures are cumulative impacts (increments to a stock) over the period of observation.

Table 4 shows estimated program impacts on capital stocks. The unadorned binary impact ITT estimates in column (1) are positive ($589), but not statistically significant. Including covariates increases the precision of the estimates, and column (2) shows estimated impacts of almost $650, significant at the 10-percent level. Column (3) again allows the program impact on capital accumulation to vary with the gender of the farmer; for this outcome the coefficient is both small and insignificant.

Columns (4) and (5) of Table 4 report the LATE estimates with and without covariates. The estimated

[11]Some elements of fixed capital were difficult to value as they were often constructed by the farmer rather than purchased on the market. RBD program staff assisted with the evaluation, but a few items (e.g. erosion barriers and certain types of fencing) are not included in our measure of fixed capital.

Table 4: Impact of RBD Program on Farm Investment: ANCOVA Estimates

	ITT			LATE		LATE (Complier sample)		
Treatment	588.9	643.8**	634.6**	906.8	971.6**	545.1*	754.3**	774.9**
	(365.7)	(287.0)	(296.3)	(557.8)	(429.5)	(306.0)	(293.9)	(324.4)
Baseline farm investment	0.94***	0.96***	0.96***	0.94***	0.96***	0.99***	0.98***	0.98***
	(0.050)	(0.027)	(0.028)	(0.050)	(0.027)	(0.032)	(0.034)	(0.034)
Education, household head		56.2	56.2		51.6		4.40	4.54
		(41.8)	(41.8)		(40.5)		(42.9)	(43.0)
Experience with target crop		15.8**	15.8**		15.9**		6.48	6.36
		(7.45)	(7.50)		(7.35)		(9.09)	(9.10)
Household size		102.6*	102.5*		92.1		89.1	89.0
		(60.0)	(59.8)		(58.6)		(77.4)	(77.7)
Land owned (manzanas)		-13.0	-13.0		-13.2		4.03	4.05
		(15.6)	(15.7)		(15.4)		(3.01)	(3.03)
Share of improved seeds		1020.5***	1021.1***		1061.7***		931.6**	931.5**
		(358.9)	(358.9)		(361.9)		(379.0)	(379.5)
Female household head		-57.6	-98.5		-71.3		23.1	112.0
		(378.5)	(333.0)		(371.3)		(477.4)	(356.7)
(Female head) x (treatment)			72.2					-160.7
			(608.0)					(741.5)
Constant	-440.4*	-1512.5***	-1507.2***	-468.7*	-1462.2***	-456.0*	-1426.3**	-1435.6**
	(236.6)	(538.8)	(535.2)	(248.4)	(522.7)	(257.6)	(703.5)	(692.8)
N	1536	1449	1449	1536	1449	1038	995	995
$\overline{R}^2$	0.861	0.880	0.880	0.861	0.881	0.867	0.894	0.894

Cluster-robust standard errors in parentheses

* p<.1, ** p<.05, *** p<.01

All regressions control for crop fixed effects. Share of improved seeds measures the share of seasons at baseline that the beneficiary household used improved seeds impacts of the program on farm investment are around \$970 and significant at the 10-percent level. At baseline, the average household in our sample had around \$7,000 in total farm capital, so these program impacts correspond to an increase, on average, of 15 percent over the baseline capital stocks. The results from the complier sample are again consistent with the standard LATE estimates.

Household Consumption

The ultimate goal of the RBD was to boost the living standards of small-scale farm families. To investigate impacts that proxy for this dimension, we adopted the household expenditure module utilized in Nicaraguan living standards surveys. As with our money-metric outcome measures, we transformed consumption expenditures into 2005 purchasing-power-parity adjusted US\$. Because the number of household members fluctuates both within and between years, we adjusted the different expenditure components by potentially different household sizes to arrive at per-capita measures. Specifically, food expenditures was converted to a per-capita measure using as a denominator the number of household members who had actually been in residence during the short recall period used to measure food spending. Other expenditure categories with longer recall periods were adjusted using the full roster of household residents (defined as those who habitually reside and sleep in the household).

As can be seen in Table 5, the effect of the program on consumption are small in magnitude, but negative for the full sample (columns (1) - (5)) and positive for the complier sample. The estimates are not statistically significantly different from zero. Again, column (3) shows results in which the impacts are allowed to differ by the gender of the RBD program beneficiary: the program effects on household consumption seem to vary significantly by gender, by a coefficient corresponding to roughly \$650. The overall treatment effect for women beneficiaries is positive but not statistically different from zero. While the gender differences for the other outcome variables are insignificant, the differential impacts of the program on expenditures by gender are intriguing. Together with the positive (albeit insignificant) effect of the program on farm income for women beneficiaries, this result may weakly suggest that women beneficiaries allocate program-induced income increases differently than do men.

3.2 Continuous treatment estimates

As discussed in Sections 1 and 2, there are multiple reasons to believe that the impacts of this type of program might evolve over time. To capture the potentially non-linear impact or duration response functions, whereby impacts depend on how much time has passed since the producer enrolled in the RBD program, we exploit the fact that treatment was rolled out in a staggered fashion within the early and late treatment groups, creating

Table 5: Impact of RBD Program on Household Consumption: ANCOVA Estimates

	ITT			LATE		LATE (Complier sample)		
Treatment	-71.3	-64.2	-147.9	-109.9	-96.9	148.0	135.9	64.8
	(123.7)	(128.4)	(130.9)	(189.1)	(191.7)	(162.0)	(172.2)	(179.2)
Baseline expenditures	0.45***	0.37***	0.37***	0.45***	0.37***	0.38***	0.29***	0.29***
	(0.056)	(0.065)	(0.065)	(0.055)	(0.064)	(0.058)	(0.059)	(0.059)
Education, household head		103.3***	103.3***		103.7***		101.3***	101.3***
		(22.5)	(22.4)		(22.6)		(29.9)	(29.9)
Experience with target crop		14.3**	14.5**		14.3**		18.9**	19.3**
		(5.74)	(5.76)		(5.66)		(7.84)	(8.00)
Household size		-156.0***	-156.7***		-155.1***		-191.8***	-192.0***
		(27.9)	(28.0)		(28.2)		(34.2)	(34.2)
Land owned (manzanas)		1.96	1.88		1.99		3.75*	3.68*
		(1.81)	(1.77)		(1.78)		(2.22)	(2.17)
Share of improved seeds		75.6	81.6		72.4		139.9	141.5
		(193.3)	(194.6)		(191.3)		(242.7)	(242.2)
Female household head		108.4	-253.6		109.8		222.9	-75.7
		(198.4)	(206.5)		(194.6)		(265.3)	(256.6)
(Female head) x (treatment)			645.8*					544.2
			(388.3)					(506.8)
Constant	1209.4***	1481.2***	1531.3***	1213.1***	1476.4***	1282.1***	1622.8***	1658.9***
	(170.7)	(281.3)	(274.7)	(172.4)	(276.1)	(201.9)	(332.0)	(327.2)
N	1579	1487	1487	1579	1487	1065	1019	1019
$\overline{R}^2$	0.402	0.454	0.455	0.400	0.453	0.330	0.388	0.389

Cluster-robust standard errors in parentheses

* p<.1, ** p<.05, *** p<.01

All regressions control for crop fixed effects. Share of improved seeds measures the share of seasons at baseline that the beneficiary household used improved seeds variation in the duration of treatment (as shown in Figure 2). The coefficient estimates of ζ_1, ζ_2, and ζ_3, from estimating equation 4, our preferred cubic specification, are shown in Table 6. The rest of this section will present these results graphically since the temporal path is somewhat hard to infer from the coefficients alone.

Table 6: Impact of RBD Program on Farm Income, Investment and Household Consumption: Fixed Effects, Continuous Treatment Estimates

	Farm Income	Investment	Consumption
Number of months in program	324.1***	219.9	-24.1
	(118.5)	(160.1)	(42.2)
Number of months in program^2	-12.2*	-2.64	1.79
	(6.90)	(11.5)	(2.05)
Number of months in program^3	0.14	-0.0084	-0.031
	(0.11)	(0.20)	(0.032)
Constant	-1644.2	-39461.8	-11018.0
	(3675.5)	(30534.8)	(6657.0)
N	3062	3147	3198
$\overline{R}^2$	0.305	0.202	0.192

Standard errors in parentheses

* p<.1, ** p<.05, *** p<.01

Not shown in table: time and crop dummies, and Mundlak instruments for fixed effects

Farm Income

Duration of time in program has a statistically significant impact on gross income in the treated activity. Drawing out the implications of the estimates shown in Table 6, Figure 4 graphs the estimated cubic relationship. Predicted farm income at the start of the program is on average roughly $7,700. Income increases over the first two years in the program, flattening out at an average predicted farm income of $10,300, for an impact estimate of roughly $2,600. Compared to the point estimate of our preferred specification from Section 3, which is around $1,400, it seems that the binary estimates may be underestimates of the fuller impacts. From the graph, it appears as though most of the benefits of the program occurred during the 24 months during which farmers were actively enrolled in the program, then flattening out. That said, incomes remain at the higher level, suggesting that a temporary intervention that offers subsidies sticks and has lasting impacts as in the Carter, Laajaj and Yang (2013) study of Mozambique.

Investment

Figure 5 plots the estimated impacts on investment, together with a 95% confidence band. As can be seen, the estimated impact of the program on beneficiaries' total capital stock increases significantly over the duration of the project, continuing to rise even after the end of active programming (24 months). The

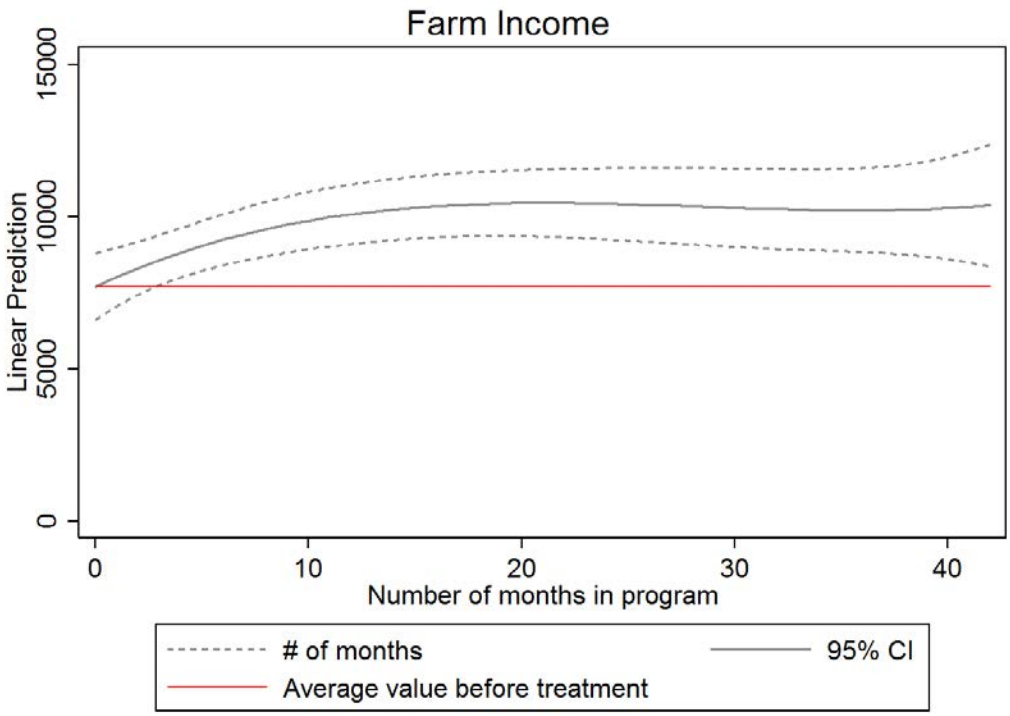

Figure 4: Predicted Farm Income by Months of Treatment

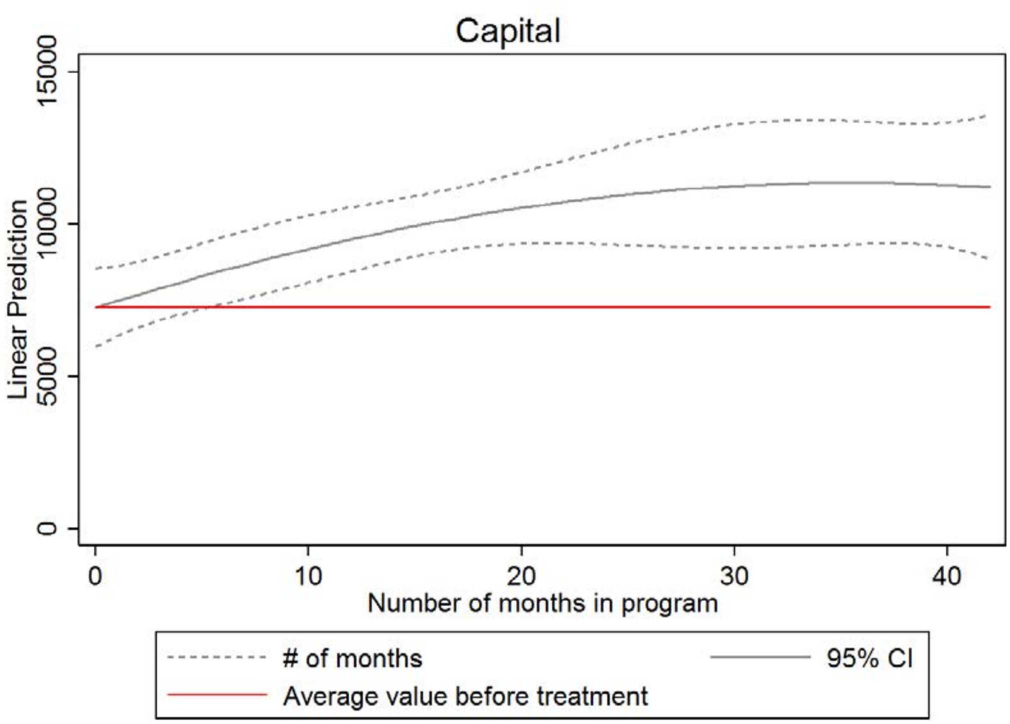

Figure 5: Predicted Investment by Months of Treatment

predicted investment level at the start of treatment is around \$7,500, and by month 42 this has risen to \$11,200 – an increase of more than \$3,700. This is well in excess of the binary LATE estimates, which suggested impacts just under \$1,000. The difference between the two types of estimates is thus greater for investment than it is for gross farm income, indicating that perhaps the increased incomes from the farm continue to be reinvested into the productive activities, rather than increasing household consumption. We therefore turn to the impacts on household consumption next.

Household Consumption

With consumption, there are reasons to suspect an initial fall in consumption since households were to self-finance program investments, which they may do by reducing consumption (especially if they are credit constrained). Some 40% of our sample is reported to be credit-constrained (in the sense of having unmet demand for loans they would like to take), suggesting that this segment of the beneficiary population may need to reduce consumption in order to self-finance up-front investments. Column 3 of Table 6 shows the estimated cubic function of treatment duration. The individual coefficients are small in magnitude and not statistically significantly different from zero. Of course, the key question is the statistical significance of the overall impact duration relationship. Figure 6 displays the cubic relationship as well as the 95% interval estimate of the duration response implied by the cubic estimates. As can be seen, the point estimates show no signs of consumption growth over the time of the program, and the interval estimator always includes zero. Given the findings of significant impacts on RBD-targeted income and on capital investment, the lack of a significant impact on living standards is somewhat puzzling. There are several potential explanations for this result, between which we cannot fully distinguish in this analysis. It could be that *total* income did not increase for beneficiary households (as opposed to income from only RBD-targeted activities). Alternatively, total income did increase, but most of it was allocated to investment rather than towards increasing immediate living standards, as indicated by the investment results in section 3.1 above.

To dig deeper into the question of total income, we examine two measure of maize production: land allocated to maize vs. the target crop, as well as income from maize. Note that the maize information is less reliable than the other survey data: the questionnaire did not elicit the maize prices received by the beneficiary households, but we base our measure on prices collected by the Nicaraguan survey team at the regional level in the year that the baseline was collected. Maize production is valued at this rough price estimate, and adjusted for inflation. As can be seen in the top panel of Figure 7, the total value of maize production declines over the study period, but it is rather imprecisely estimated. The bottom of Figure 7 shows the land planted to corn and to the target crop, respectively. It appears that some substitution indeed took place, as area planted in corn declines while area dedicated to the target activity increases. This

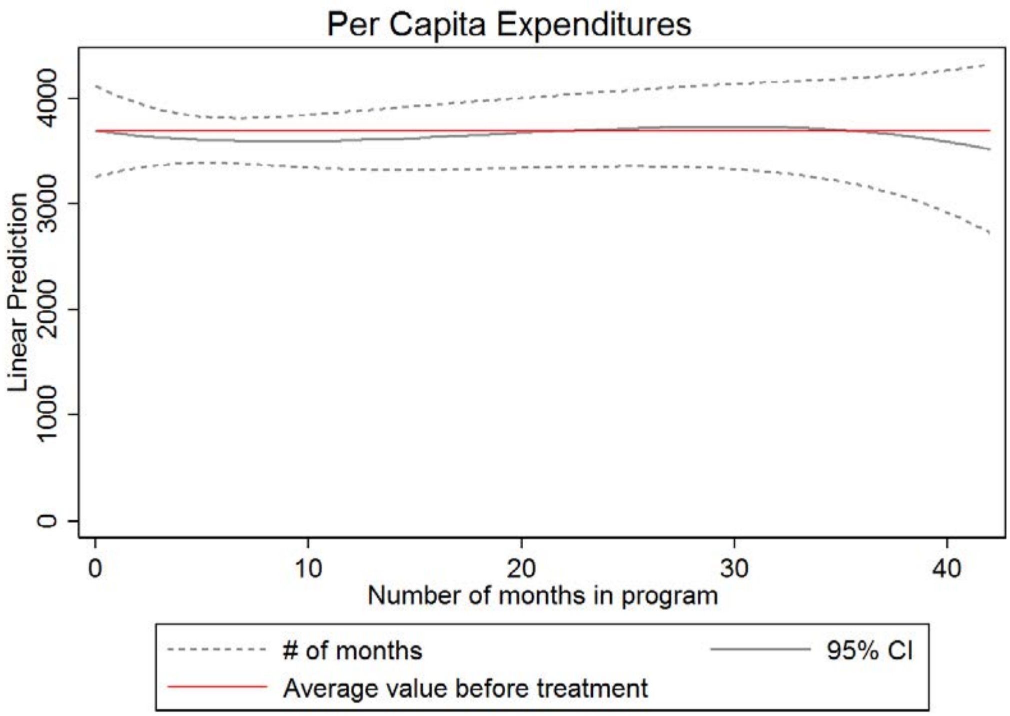

Figure 6: Predicted Household Consumption by Months of Treatment

suggests that the total impact on household income is lower than that estimated above. However, since we do not have good information on what happened to maize prices during this time, these results should be taken with a grain of salt. The fact that households increased their on-farm investment still implies that the beneficiaries of the program felt that the increased income from the target activity was worth investing in.

Finally, it is also possible that impacts are quite heterogeneous and that average effects of the sort considered in this section disguise the impacts of a business development program in a world where not everyone succeeds as a small-scale agricultural entrepreneur. The next section therefore looks more carefully at program impacts at different parts of the distribution.

4 Impact Heterogeneity

There are multiple reasons why programs like the RBD may have heterogeneous impacts, including:

1. *Heterogeneous access to the financial capital needed to make the most of the RBD intervention;*

2. *Complementarity between the RBD intervention and unobservable assets that are not equally distributed across the population, such as farming skills, learning capacity and business acumen; and,*

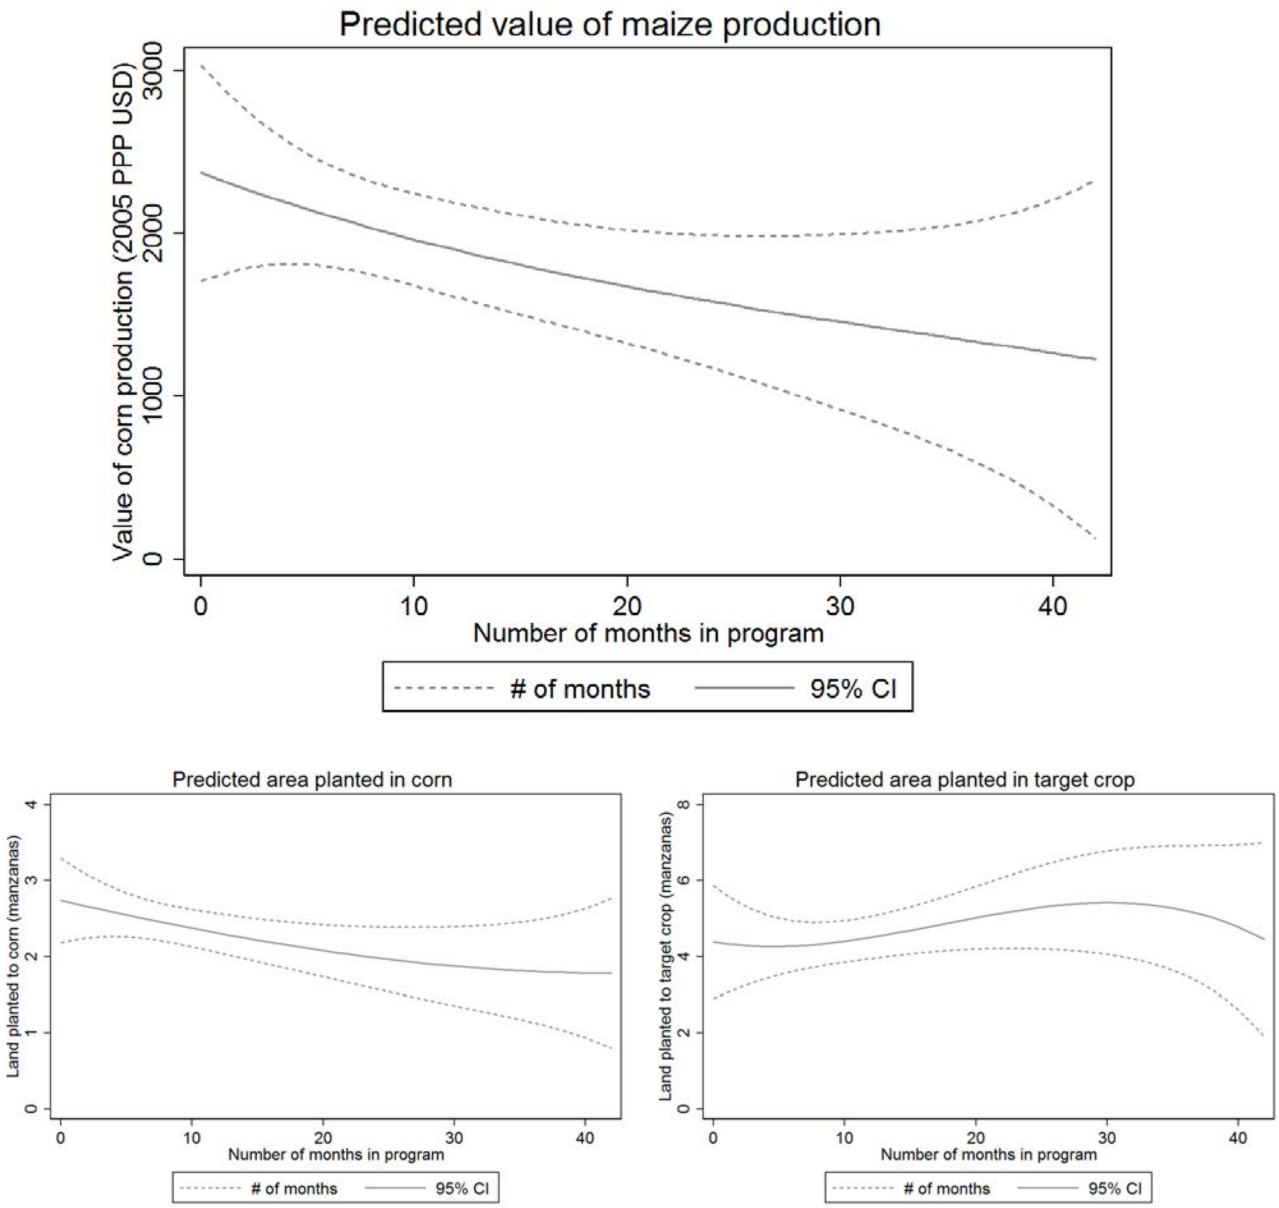

Figure 7: Predicted Maize Production, Area Planted in Corn, and Area PLanted in Target Crop, by Months of Treatment

3. *Differential luck, with some succeeding and others failing for stochastic reasons.*

Earlier analysis conducted with only the mid-line data revealed substantial evidence of impact heterogeneity, with the program showing few impacts on the well-being of the poorest-performing 50% of the population (when compared against the poorest-performing segment of the untreated households), with quite high returns to the best performing segment of the treated group, when compared against top performers in the then untreated control group (Toledo and Carter (2010)). In this section, we use all three rounds of data and our continuous treatment model to further explore impact heterogeneity.

4.1 Econometric Approach

Conventional regression methods (such as those just employed above in Section 3) estimate average or mean relationships. They assume that the vector of covariates affects only the *location* of the conditional distribution of y, not other aspects of y's conditional distribution. Conditional quantile regression methods allow us to see whether the statistically average relationship is in fact a good description of the relationship in all parts of the distribution. Specifically, quantile regression allows us to recover the regression parameters that best describe the impacts on observations in different portions of the error distribution for our regression model.

Observations in the higher quantiles are those that "do better" than would be predicted by the household's level of treatment and other regression variables (*e.g.*, are in the upper tail of the conditional per capita consumption distribution). For simplicity, we will refer to these observations in the higher quantiles as "high performers," although absent other evidence, this should be taken to mean high-performing observations and not necessarily high-performing household types. Conversely, observations in the lower quantiles are those are in the lower tail of the conditional distribution of the outcome variable. Quantile regression allows us to see if the marginal impact of RBD program participation at various parts of the conditional distribution of the outcome variables differs from the impacts at the mean—*i.e.* the average relationship estimated in Section 3. Note that if the average regression model explains the data well, the impact estimates (the slope of the impact duration function) will be the same for all quantiles. However, if there is unobserved heterogeneity in the impacts, then the impact slopes across quantiles may be different. As mentioned above, there are conceptual reasons to suspect that the RBD program might have heterogeneous impacts. Not that reason 3 above for heterogeneity would imply "high-performing observations," whereas reasons 1 and 2 would imply the existence of high-performing households.[12]

[12] Note that if heterogeneity is driven by capital constraints (reason 1), then low performance could be attenuated by augmenting business services with a credit program. However, if low performers lack fundamental human capital, it is less obvious how to ameliorate low performance.

To recover conditional quantile estimates, we employ the method developed by Abrevaya and Dahl (2008) that extends a correlated random-effects framework (like regression model (3) above) to apply to conditional quantile models. While quantile models have been widely used in empirical studies since their development by Koenker and Bassett (1978), they are not often applied to panel data, likely because of the difficulty of differencing in the context of conditional quantiles. This problem arises because quantiles are not linear operators, so that, simply put, the conditional quantile of a difference is not simply a difference of the conditional quantiles. Importantly, this methodology based on correlated random-effects preserves the fixed effects characteristics of the results, inoculating them against systematic or spurious correlation between the duration of treatment and initial and time-invariant conditions. Note also that the conditional errors are estimates of $\upsilon_i + \varepsilon_{it}$ from equation 3. That is, the error contains the time-invariant, random effect component.

4.2 Generalized Quantile Estimates

This section explores the heterogeneity of the impact or duration response function by estimating the conditional quantile functions for our preferred (cubic) parametric continuous treatment models. Parameter estimates for the Abrevaya and Dahl (2008) estimator can be obtained with any quantile regression package. Standard errors are obtained through bootstrapping (we use 500 replications), drawing households with replacement from the sample and estimating the variance-covariance matrix from the resulting empirical variance matrix. We present the results graphically, showing the predicted values of the outcome variables as a function of the length of time in the program, for the 25th, median and and 75th quantiles, with bootstrapped 95% confidence intervals displayed as dotted lines around the point estimates. The regression coefficients can be found in Appendix table 9.

Figure 8 (a) displays the results from the quantile analysis of income in the targeted activity. As can be seen, these estimates confirm the hypothesis that program impacts are heterogeneous across the participant population. The impacts of the program are greatest at the high end of the distribution, with the 25th and median impacts still significantly different from zero, but substantially smaller in magnitude. The high performers in the 75th quantile experience a steeper impact response function that also peaks later than the lower quantiles. Indeed, towards the end of the program duration, despite dropping down a little from the peak, farm incomes are roughly \$1,500 greater than at the program beginning, more than three times the long-term impact level for the producer at the median or 25th quantile of the conditional income distribution, both of whom increase their incomes by slightly more than \$400. The statistical significance of these impact paths can be approximated by comparing the confidence interval to the red line, which is drawn at the

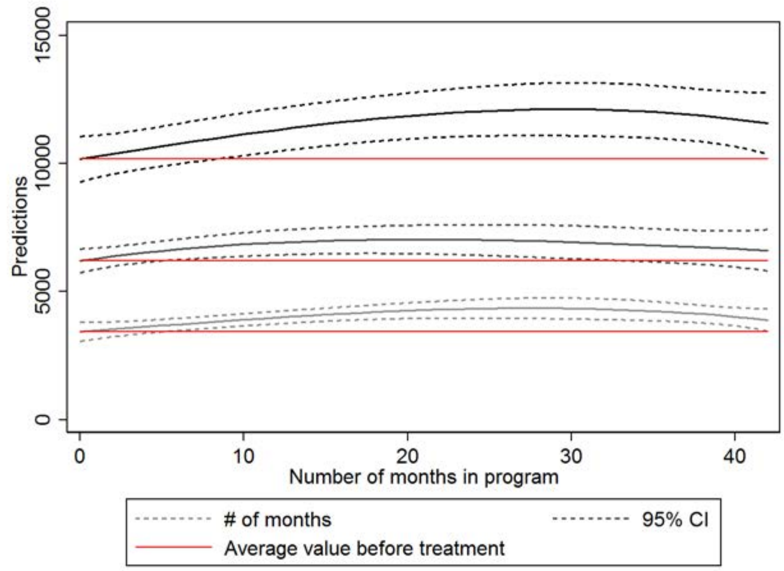

(a) Farm income

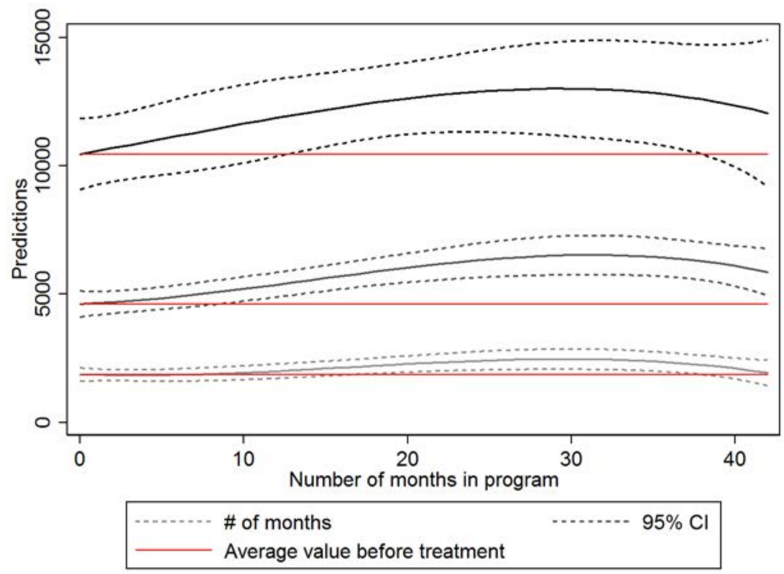

(b) Investment

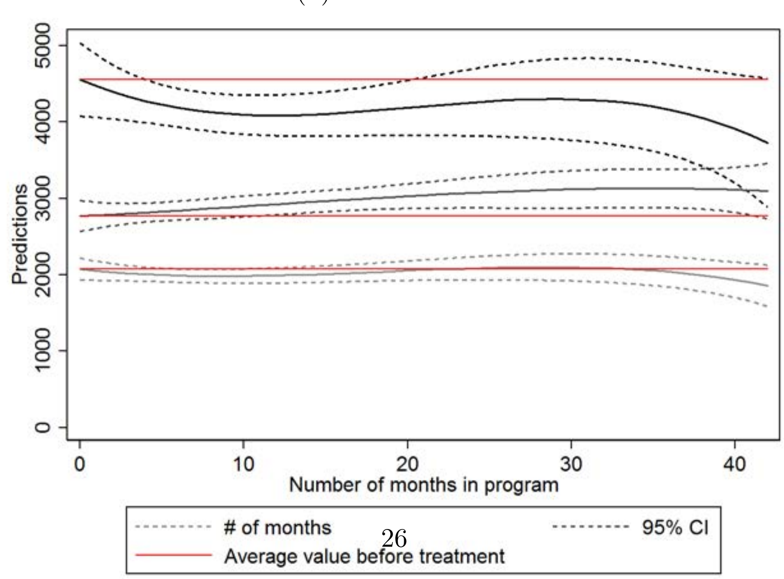

(c) Consumption

income level at zero months of treatment.

The estimated impacts of the RBD program on capital also vary substantially across conditional quantiles (Figure 8 (b)). The impact on investment increases as we move upwards in the conditional distribution of capital. For households in the 75th quantile, investment increases by roughly \$1,600 over the course of the program, a magnitude very similar to the farm income increases seen in the panel above. At the 50th conditional quantile, households also increase their investment by a substantial amount: roughly \$1,200. The lowest quantile, however, sees the amount of investment in capital dip back down towards the end of the program, basically returning to baseline levels.

For per capita consumption, we see a different pattern, whereby the median regression shows significant positive impacts on consumption of around \$400, although the confidence intervals get a bit wider at the very end of the duration. Households at the 25th and 75th conditional quantiles experience no impacts on consumption; in fact, for the 75th quantile, the impacts on consumption are in fact negative for much of the period.

4.3 Interpreting Impact Heterogeneity

It is tempting to interpret this impact heterogeneity as signaling that the RBD program did not work for everyone. However, as discussed above, it is possible that the lower quantiles are comprised of observations in which output was diminished by a negative shock. For example, a program like the RBD would be unlikely to have any impacts in the face of a localized drought since improved varieties, marketing channels, *etc.* would be useless if production were reduced to zero by weather events.

One way to gain purchase on this problem and to garner some insight on the source of this heterogeneity is to ask whether the same households consistently occupy the same quantile position in the conditional error distribution. If they do–meaning there are consistently upper quantile households and consistently lower quantile household–then we have evidence that program impacts vary systematically by (unobserved) household type.

To explore this idea, we recovered the residual for each observation in each round from a median regression, and denote by q_{it} the quantile in which household i's residual was out of the distribution of errors in round t. Using a standard analysis of variance decomposition, we can decompose the total variation in q_{it}as follows:

$$\sum_{i=1}^{N}\sum_{t=1}^{3}(q_{it}-\bar{q})^2 = \sum_{i=1}^{N}\sum_{t=1}^{3}(q_{it}-\bar{q}_i)^2 + 3\sum_{i=1}^{N}(\bar{q}_i-\bar{q})^2,$$

where N is the number of households in the dataset, $\bar{q}$ is the overall mean in the dataset, while $\bar{q}_i$ is the mean quantile for household i over the 3 rounds of the data. The first term on the right hand side is the within sum of squares (WSS), while the second term is the between sum of squares (BSS). Note that if no household changed position in the error distribution from year to year, then the WSS would be zero. Conversely, if a household's error quantile varied randomly from year to year (sometimes high, sometimes low and sometime in between), then $\bar{q}_i \approx \bar{q}\,\forall i$ and the BSS would be a small fraction of the total variation in q_{it}.

Table 7 shows that the fraction of total variation that is between households ranges from 66% to 83% for our three primary outcome variables. While a modest fraction of the variation comes from households moving between error quantiles over time, the bulk of the overall variation is coming from time invariant differences between households. In other words, there is evidence that particular households tend to occupy upper quantiles, and others tend to occupy lower households. Given that the latter seem to enjoy little benefit from the RBD program across all three indicators, this finding suggests that the RBD program is a blunt instrument for improving the economic welfare of an important subset of household, but quite effective for those who have the right complementary market access and, or skills.[13]

Table 7: Within vs. Between Variation in Error Percentiles

	Within variation (%)	**Between variation (%)**
Income	33.6	66.4
Total Investment	17.4	82.6
Consumption	27.3	72.7

5 Conclusion

Nicaragua's Rural Business Development Program (RBD) was a 24 month intervention designed to boost the productivity and incomes of a largely poor, rural farming population by enhancing their business knowledge and improving their access to markets and technologies. Gauging the effectiveness of a program like the RBD that addresses the rural poor as incipient entrepreneurs rather than as passive recipients of transfers, faces several challenges. The first challenge is to gauge the long-term impacts of the RBD intervention. Beneficiaries may continue to learn, invest and realize further benefit from the intervention long after the 24 month period. It is of course also possible that any immediate impacts are not sustained if beneficiaries return to their prior status after the period of direct intervention ends. The second challenge is to understand the

[13]Efforts reported in Toledo (2011) to unpack the reasons behind the impact heterogeneity reported in Toledo and Carter (2010) are only partially satisfying. That analysis focused on explanation (1) above, categorizing households based on their credit-rationing status. While credit market status is of course endogenous, that analysis revealed no simple relationship between performance and contemporaneous credit rationing status. Indeed, the only factor uncovered was past credit history. RBD impacts on farms with prior credit history appeared quite large and significant. Unfortunately, the interpretation of prior credit history as a factor explaining heterogeneous program impacts is ambiguous. It seems most likely that those with past credit histories are actually those with higher levels of farming and business acumen (pointing toward explanation 2 above). It may also be that those acumen levels were themselves endogenously produced by prior random or prior program-based access to credit (and business opportunities).

heterogeneity of impact across the target population, both in terms of participation in a somewhat complex program, and in terms of the business success of those who choose to participate.

To address these challenges, we employed a 5-year roll-out design that randomized beneficiaries' exposure to the RBD program. The design also allowed us to identify a two-sided complier sample and to focus the analysis only on those who (eventually) enrolled in the program. Using 3 rounds of data from this design, we find that on average has substantial impacts on income in the targeted activities ($1,400 to $2,600 annually)[14] and on agricultural investment ($3,700). Estimated impacts reach their maximum approximately 24 to 30 months after initial exposure to the program. Somewhat surprisingly, there are virtually no impacts on household consumption expenditures,[15] and in fact some weak evidence that the program *reduced* expenditures, as would be expected if beneficiaries were investing more in their farm but facing liquidity constraints.

At a direct program cost of $2,500 per-farmer enrolled in the program, these average estimates indicate that the RBD was a cost-effective instrument for boosting the average income and assets of its beneficiary farmers. However, its effectiveness as an instrument to address rural poverty depends on the distribution of impacts across the program's overall target population. Looking at the full distribution of impacts is especially important for efforts like the RBD program that target beneficiaries' income-generating and entrepreneurial capacities.

In the first instance, we note that just over one-third of the target population declined to participate in the program. The one-third who did not participate had modestly lower living standards at baseline.[16] In addition, we employed fixed-effects analogue conditional quantile regression methods to explore the degree to which the average pattern of impact faithfully reflects the full distribution. We find evidence of significant heterogeneity in impacts on program income and investment, with smaller program impacts in the lower quantiles of the conditional error distribution. Observations in the 25th quantile show long-term income and investment impacts that are one-third to one-half the size of the estimated average impacts, but still significant and positive. We further provide evidence that households' position in the error distribution is relatively constant over time, implying that lower-quantile observations are comprised of "low-performer" type households for whom RBD program impacts are less effective.

The existence of these two groups (those that did not participate, and those that experienced more modest impacts even when particpating) serves as a useful reminder that not all small farms can upgrade and succeed.

[14] As discussed above, these estimates are upper bound estimates on the impacts on total family income.

[15] This insignificant impact is not a "precisely estimated zero." Indeed, the point estimate of the long-term impact is economically large, but its standard error is large, potentially indicating significant heterogeneity in program impacts on consumption.

[16] Recent work by Macours and Vakis (2008) and Laajaj (2012) on poverty and aspirations suggest that there may be some individuals who could benefit from interventions such as the RBD, but that they need smaller, confidence, and aspiration-building steps before they are willing to jump into a more forward-looking and entrepreneurial profile.

If the goal is to eliminate rural poverty, then this limitation needs to be kept in mind as other interventions may be needed to improve prospects for this sub-population and their children. Looking forward, it may be that next-generation RBD programs can reduce the size of this minority. While the analysis here is unable to identify which families failed to succeed and why,[17] it is likely that some failures were due to the natural vagaries of agriculture as a risky activity. Efforts to incorporate elements of insurance into small farm development strategies may have a key role to play in this regard, allowing a greater percentage of the small farm population to succeed over the longer term.

In addition, the RBD program did not include a direct credit market intervention. The overall MCC program in Nicaragua operated in part on the theory that improved property registration would indirectly improve smallholder access to capital by increasing their collateral and credit-worthiness to the extant banking sector. Whether or not that strategy would have worked remains an open question, as the property registration component of the program was eliminated in early 2009 (see footnote 5 above). The observed pattern of increasing income but sluggish changes in living standards might signal the existence of capital constraints, with income increases soaked up to self-finance future fixed and working-capital investments.

[17]One important message that emerged from the mid-line evaluation is that there is no evidence that farms closer to the asset minima benefited less from the program than did better endowed farmers (Carter and Toledo, 2011). While the asset floors and ceilings used to establish RBD eligibility were based on best-practice intuition, it is clear from a targeting perspective that more work needs to be done to see if there is such a thing as a farm that is too small to benefit from this kind of intervention.

References

Abrevaya, Jason, and Christian M Dahl. 2008. "The Effects of Birth Inputs on Birthweight." *Journal of Business & Economic Statistics*, 26(4): 379–397.

Aguero, J., M. R Carter, and I. Woolard. 2010. "The impact of unconditional cash transfers on nutrition: The South African Child Support Grant." Working Paper.

Anderson, J., and G. Feder. 2003. "Rural extension services." *World Bank Policy Research Working Paper No. 2976*.

Ashraf, Nava, Xavier Giné, and Dean Karlan. 2009. "Finding Missing Markets (and a Disturbing Epilogue): Evidence from an Export Crop Adoption and Marketing Intervention in Kenya." *American Journal of Agricultural Economics*, 91(4): 973–990.

Banerjee, Abhijit, Esther Duflo, Raghabendra Chattopadhyay, and Jeremy Shapiro. 2011. "Targeting the hard-core poor: an impact assessment."

Bardhan, Pranab, and Dilip Mookherjee. 2011. "Subsidized Farm Input Programs and Agricultural Performance: A Farm-Level Analysis of West Bengal's Green Revolution, 1982 - 1995." *American Economic Journal: Applied Economics*, 3(4): 186–214.

Buchinsky, M. 1998. "Recent advances in quantile regression models: a practical guideline for empirical research." *Journal of Human Resources*, 88–126.

Carter, Michael R., Rachid Laajaj, and Dean Yang. 2013. "The Impact of Voucher Coupons on the Uptake of Fertilizer and Improved Seeds: Evidence from a Randomized Trial in Mozambique." *American Journal of Agricultural Economics*.

Carter, Michael R., Rachid Laajaj, and Dean Yang. 2014. "Subsidies and the Persistence of Technology Adoption: Field Experimental Evidence from Mozambique." National Bureau of Economic Research Working Paper 20465.

Chamberlain, G. 1982. "Multivariate regression models for panel data." *Journal of Econometrics*, 18(1): 5–46.

Chamberlain, G. 1984. "Panel Data. Handbook of Econometrics." *Griliches and M. Intriligator, eds.*

Cole, Shawn Allen, and A. Nilesh Fernando. 2012. "The Value of Advice: Evidence from Mobile Phone-Based Agricultural Extension." Social Science Research Network Harvard Business School Finance Working Paper 13-047, Rochester, NY.

Evenson, Robert E. 2001. "Chapter 11 Economic impacts of agricultural research and extension." In *Agricultural Production.* Vol. Volume 1, Part A, 573–628. Elsevier.

Feder, Gershon, Roger H. Slade, and Lawrence J. Lau. 1987. "Does Agricultural Extension Pay? The Training and Visit System in Northwest India." *American Journal of Agricultural Economics*, 69(3): 677–686.

Hanlon, Joseph, Armando Barrietnos, and David Hulme. 2010. *Just Give the Poor the Money: The Development Revolution from the Global South.* Kumarian Press.

Hirano, Keisuke, and Guido W. Imbens. 2004. "The Propensity Score with Continuous Treatments." In *Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives.* , ed. Andrew Gelman and Xiao-Li Meng, 73–84. John Wiley & Sons, Ltd.

Keswell, M., and M. Carter. 2011. "Poverty and Land Distribution."

Keswell, M., and M. R. Carter. 2014. "Poverty and Land Distribution." *Journal of Development Economics*, 110: 250–61.

King, Elizabeth M., and Jere R. Behrman. 2009. "Timing and Duration of Exposure in Evaluations of Social Programs." *The World Bank Research Observer*, 24(1): 55–82.

Koenker, R., and G. Bassett. 1978. "Regression quantiles." *Econometrica: journal of the Econometric Society*, 33–50.

Laajaj, R. 2012. "Closing the Eyes on a Gloomy Future: Psychological Causes and Economic Consequences." Working Paper.

Macours, K., and R. Vakis. 2008. "Changing households' investments and aspirations through social interactions: Evidence from a randomized transfer program in a low-income country." *World Bank Research Paper. World Bank, Washington, DC.*

McKenzie, David. 2012. "Beyond baseline and follow-up: The case for more T in experiments." *Journal of Development Economics*, 99(2): 210–221.

Mundlak, Y. 1978. "On the pooling of time series and cross section data." *Econometrica: journal of the Econometric Society*, 69–85.

Tjernström, Emilia, Michael R. Carter, and Patricia Toledo. 2013. "Identifying the Impact Dynamics of a Small Farmer Development Scheme in Nicaragua." *American Journal of Agricultural Economics, Papers and Proceedings*, 95(5): 1359–1365.

Toledo, Patricia. 2011. "Impact Evaluation of a Rural Business Program Using Field Experiment Data." Manuscript.

Toledo, Patricia, and Michael R. Carter. 2010. "Impact of Business Services on the Economic Well-being of Small Farmers in Nicaragua."

World Bank, the. 2008. "Nicaragua - Poverty Assessment (Vol I: Main report)." 39736, Washington, D.C.

A Appendix A - Eligibility Criteria by Productive Activity

Table 8: Eligibility Criteria Used to Identify Farmers in Target Activities

	SESAME	BEANS	VEGETABLES	CASSAVA	LIVESTOCK
Asset Floor*	7 hectares	3.5 hectares	1.4 hectares	3.5 hectares	10 mature cows
Asset Ceiling	35.2 hectares	35.2 hectares	14.1 hectares	70.4 hectares	100 mature cows
Prior Experience	1.4 hectares in sesame	0.7 hectares in beans	Some vegetable production	1.4 hectares in cassava	Developed livestock activity
Water	−−	−−	On-farm water source	−−	On-farm water source
Legal Status	Farmer has land title or is in possession of land				
Age	Farmer must be at least 20 years old				
Environment	Land located outside of national protected areas				

*Minimum farm size reduced when farm is irrigated

B Appendix C - Quantile regression tables

Table 9: QUANTILE REGRESSION COEFFICIENTS

	Income			Investment			Consumption		
Quantile:	25^{th}	50^{th}	75^{th}	25^{th}	50^{th}	75^{th}	25^{th}	50^{th}	75^{th}
Months treated	48.3	91.6*	103.1	-20.4	27.7	113.7	-23.2	10.5	-90.6*
	(34.3)	(48.9)	(80.7)	(35.4)	(50.0)	(159.3)	(14.6)	(24.7)	(50.0)
Months treated^2	0.13	-2.98	-0.28	3.43	4.05	1.15	1.73**	0.30	5.35*
	(2.17)	(2.88)	(4.65)	(2.31)	(3.40)	(10.9)	(0.86)	(1.54)	(2.94)
Months treated^3	-0.024	0.025	-0.033	-0.069*	-0.095	-0.070	-0.031**	-0.0087	-0.087*
	(0.037)	(0.047)	(0.076)	(0.039)	(0.059)	(0.18)	(0.014)	(0.026)	(0.049)
Farmer age	-5.53*	-3.17	-2.67	14.8***	27.7***	63.5***	9.63***	15.3***	34.9***
	(3.34)	(5.24)	(10.3)	(3.88)	(8.94)	(20.7)	(2.60)	(3.61)	(6.02)
Farmer education	19.1	77.7***	174.1***	103.8***	315.5***	655.2***	115.2***	176.6***	263.8***
	(15.4)	(25.2)	(56.2)	(24.6)	(57.9)	(115.9)	(13.7)	(15.9)	(26.6)
Constant	-8722.1	-5427.6	-5689.8	-9598.7	-35253.5	-50763.3	-3403.0	-8388.2	-26102.4*
	(7968.5)	(6606.1)	(7819.3)	(19406.3)	(36463.6)	(73833.4)	(4790.3)	(9471.0)	(13684.2)
N	3062	3062	3062	3147	3147	3147	3198	3198	3198
Pseudo-R^2	0.21	0.30	0.33	0.078	0.14	0.19	0.11	0.13	0.15

Bootstrapped standard errors in parentheses, from 500 reps drawn with replacement over households

The regressions also include time fixed effects, activity fixed effects and averages of the time-varying variables

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

C Appendix C - Hypothesis Testing

To test whether the differences that we observe between the different quantiles are statistically significant, we employ the minimum-distance framework in Abrevaya and Dahl (extended from Buchinsky's (1998) framework to the panel data context) to test the equality of the parametric duration response variables' effects across quantiles. Since both *months*, $months^2$, and $months^3$ enter into our preferred cubic model, the relevant test is a joint test of equality. In other words, the null hypothesis is

$$H_0: \ \zeta_{1,\tau_1} = \zeta_{1,\tau_2} = \zeta_{1,\tau_3} \wedge \zeta_{2,\tau_1} = \zeta_{2,\tau_2} = \zeta_{2,\tau_3} \wedge \zeta_{3,\tau_1} = \zeta_{3,\tau_2} = \zeta_{3,\tau_3},$$

where ζ_1,ζ_2 and ζ_3 are the estimated coefficients on *months*, $months^2$, and $months^3$, respectively, and τ_1, τ_2 and ,τ_3 are the different estimated quantiles ($25^{th}, 50^{th}$ and 75^{th}).

In following the Abrevaya and Dahl (2008) testing framework, and the only changes we make are to allow for an additional round of data and the fact that we include averages of the time-varying regressors, instead of their value in each round. The minimum-distance test statistic has a limiting chi-square distribution, with degrees of freedom equal to the number of restrictions (in our case 6). These test statistics and their associated p-values are shown in Table 10. As the table shows, the estimated effect of months in the program on capital vary significantly across the quantiles, but the evidence for income and consumption show weak and no difference, respectively.

Table 10: Tests of marginal-effect equality across quantiles

Outcome variable	χ^2_6-statistic	p-value
Income	9.22	0.16
Total capital	16.85	0.0098
Per-capita consumption	5.42	0.49

For each outcome variable, the p-values reported are for the null hypothesis of joint equality of the marginal effects of the variables *months*, $months^2$, and $months^3$ for the quantiles .25, .50 and .75. Results are based on 300 bootstrap replications