TESTING ALTERNATIVE APPROACHES TO CONDITIONAL CASH TRANSFER PROGRAMS IN EDUCATION: EVIDENCE FROM COLOMBIA[1]

First Draft: May 2006

Current Draft: July 2007

Preliminary and Incomplete:

Please Do Not Cite

Felipe Barrera-Osorio

(World Bank)

Marianne Bertrand

(Chicago GSB)

Leigh L. Linden

(Columbia University)

Francisco Perez

(Ministry of Education, Colombia)

[1] An undertaking of this magnitude requires the assistance of many individuals. We are most indebted to the Secretary of Education of Bogota for cooperating with us in this novel experiment, putting up with the constraints created by the research effort, and, of course, financially supporting the entire project. While everyone at the SED has been extremely helpful we are particularly indebted to Abel Rodriguez, Catalina Velasco and Margarita Vega. We are indebted to Silvia Restrepo of Fedesarrollo for there logistical assistance particular during the data collection efforts. Camilo Dominguez has done an excellent job as a research assistant during the entire project, and we thank Carlos Ospino and Lucas Higuera for their help at key points in the effort. We thank Sendhil Mullainathan and Mario Sanchez for their comments and assistance, and thank the participants at the BBL of HDN Education for their helpful question and comments. All errors are of course (and unfortunately) our responsibility. Please send correspondence to Leigh Linden at leigh.linden@columbia.edu.

Abstract:

We evaluate multiple variants of a commonly used intervention to boost education in developing countries, the conditional cash transfer (CCT). Specifically, we test three treatments: a basic CCT treatment based on school attendance, a savings treatment that postpone a bulk of the cash transfer due to good attendance to just before children have to reenroll in school, and a tertiary treatment where some of the transfers are conditional on students' graduation rather than attendance. The results suggest that changing the timing of the transfer through the savings treatment does not change families' behavior, but subsidizing access to higher education increases attendance during secondary school (by 4 percentage points) and substantially increases participation in higher education (by 49 percentage points). Our strategy also allows us to assess intra-household variations in treatment. We find that siblings of children receiving the tertiary treatment work more hours (1.7 per week) while treated children work less (3 per week).

I. Introduction

Education plays an important role in the development process. At both the macro (for example, Krueger and Lindahl, 2001) and micro level (Angrist and Krueger, 1991; Duflo, 2001, among others), there is strong evidence that education generates higher levels of both income and growth. As a result, developing countries could contribute substantially to future income growth by increasing attendance rates. The challenge, however, is getting the kids in school. For example, the net enrollment rate in primary education in 2004 in Sub-Saharan Africa, Oceania and Western Asia was 64, 80 and 83 percent respectively. Problems are more pronounced in rural areas, and in historically disadvantaged groups like girls and low-income families (United Nations, 2006).

Despite the importance of education, we are still far from understanding what determines whether or for how long children are educated. The classic model postulates a simple comparison of the future returns of additional schooling to the short-term direct costs of enrollment and the opportunity costs of the time required to attend. And while it is clear that even this simple relationship is difficult to estimate rigorously, more recent models suggest that family dynamics, peer influences, liquidity constraints or even personal commitment issues can influence the education decision process among children and their parents.

Over the last decade a large and growing literature has begun to grapple with these issues using natural and actual experiments. For example, since acquiring knowledge is the main objective to spending time in school, one would expect that students should respond to the quality of education, especially in lower income countries where the quality of education is substantially lower (Pritchett, 2004). However, improving quality does not seem to be a major inducement since interventions proven to improve the quality of education do not seem to increase participation levels much (Banerjee, Cole, Duflo, and Linden, 2007; He, Linden, MacLeod, 2007; Muralidharan and Sundararaman, 2006).

On the other hand, the short-term, direct costs and benefits of school participation do seem to have an effect. Families respond to the direct costs of enrollment by increasing enrollment when school fees are reduced (Barrera, Linden, Urquiola, 2007).

Similarly, families respond to direct inducements to attend such as meals or direct cash incentives (Vermeersch and Kremer, 2005; Schultz 2004).

Our strategy is to build on these new research findings by testing multiple variants of a single well-established program, the conditional cash transfer. By working with a large municipality, we randomly assign multiple treatments using an oversubscription model that allows us to generate exogenous variation in treatment within family, school, and peer networks. The strategy allows us to manipulate the timing of the payments families receive under the program as well as the specific incentives they face with respect to the decisions to attend and enroll in school. More generally, this strategy demonstrates a model for aligning policy makers interests in designing the most effective transfer system and the academic interest in understanding the critical system of decisions that determine the development of human capital.

The basic conditional cash transfer model was first explored in Mexico's PROGRESSA (now OPPORTUNIDADES) program. In the program, students' families received a cash transfer if they enrolled in school and attended for at least 80 percent of the days in a given month. It proved effective (Schutlz, 2004), and has expanded rapidly with at least 20 known countries conditioning transfers on either student enrollment or attendance rates. The number of evaluations of the model has grown at a similar rate, and most of these suggest that this basic intervention can increase school participation from 2-10 percent.[2] However, almost all of these studies evaluate the same basic model in which families are paid a direct subsidy for either enrolling or attending school sufficiently often.

We build on this basic model by evaluating three separate interventions in two separate experiments. First, we use the basic treatment implemented in a manner very similar to the original PROGRESSA program. This is combined with a second intervention that, using the same conditions, varies the timing with which the funds are distributed to families, distributing 2/3 of the funds to families immediately and the remaining funds at the time the students enroll in school. This treatment is designed to assess how serious savings constraints (either due to the costs of saving, individual

[2] See, among others, Attanasio et al. (2005); Behrman, Sengupta, and Todd (2005); Cardoso and Souza (2004); Chaudhury and Parajuli (2006); Filmer and Schady (2006); Glewwe and Olinto (2006); Maluccio and Flores (2005); Pitt, Khandker, and Fuwa (2003); Schady and Araujo (2006); Schultz (2004).

hyperbolic discounting, or even family commitment issue) are in determining students' enrollment and attendance patterns. Second, we test, in a second experiment, a treatment that provides children with the same lower monthly subsidy as the savings treatment, but also pays a large subsidy that incentives both graduation and matriculation to an institution of higher education.

To allocate these treatments, we use an over-subscription model rather than the basic geographic allocation strategy used in previous studies. We staged a large recruitment drive in two urban localities and, in two official public events, randomly allocated about 10,000 treatments to 17,309 registered children. This model allows us to randomize at the child-level, generating variation within schools, families, and networks of friends. By pairing this randomization with detailed information on children's siblings and friends, we are able to disentangle how these opportunities change the allocation of work in the household and the activities of the recipients' peers.[3]

Finally, we also collect attendance data through a series of school visits in order to assess the importance of self-reporting bias in the survey data used in other studies of conditional cash transfer models. This bias is particularly important in such contexts. While subjects' responses on the surveys always have no implications for their participation in the program, the subjects have already been conditioned to value attendance by the program and understand that their receipt of the transfers is determined by their rates of attendance. This could lead to a general upward bias in the reporting of attendance and could also lead to a differential bias by those most involved with the program – the treatment families.

The results suggest that all of these factors are important. All of the treatments generate significant changes in the behavior of the children and the families, increasing academic participation by 2.9 percent and increasing the quality and quantity of meals. However, the variations in the incentives did matter. The reduction in the short-term payments cause no reduction in attendance suggesting that short-term liquidity issues are less severe than previously thought. Second, changing the type of incentive also has an effect. The savings treatment encouraged slightly higher attendance rates and 4.6 percent high matriculation rate to tertiary institutions. The tertiary treatment generates

[3] Results for peer effects will appear in a future draft.

significantly higher attendance rates, increases enrollment among low attending students (1 percentage point) and generates a 48 percentage point increase in the number of children pursuing higher education. Within families, however, the tertiary benefit also seems to generate a reallocation of the workload within recipients' families with recipients working less and their school aged siblings working more.

The paper is organized as follows. First, we describe the educational system in Bogota, Colombia in the following section. In Section 3, we describe the research design, including the design of the individual treatments, the allocation process, the various data sets, and the statistical models involved in the process. We present the results of the analytical models in Section 4. Finally, we conclude in Section 5.

II. Education in Bogota

Colombia is a relatively typical middle income, Latin American country. Compared to poorer countries, child mortality is relatively low at 21 per 1000 births and individuals can expect to live long lives -- life expectance at birth is 72.6 years. The per capita income of Colombia is US$ 2,020, and 17.8 percent of the population living on less than two dollars per day (World Bank, 2006).

While the central government maintains control of curriculum, the allocation of teachers, and their wages, municipalities are primarily responsible for the administration of public education using national funds. The central government provides resources, primarily from income and VAT taxes, and 90 percent of these funds are required by law to go toward health and education. Municipalities that have greater capacity to collect and administer taxes supplement central resources with local resources, usually from property taxes. With these funds, municipalities must develop, maintain, and run the facilities in their jurisdictions.

The academic year runs from the end of January until the middle of November. The system is divided into three categories: basic primary (grades one through five), basic secondary (grades six through nine) and middle secondary (grades ten and eleven). After finishing the eleventh grade, children can matriculate to either traditional universities or one of many vocational schools. Students usually start school at five to seven years of age, and legally children are required to attend school through the ninth grade, a period referred to as basic education.

Like in most urban areas in middle-income countries, school attendance is highest for younger children. The enrollment rate for students of age between 5 and 13 are close to 100 percent. After 13 years the attendance rate starts to decline. The average attendance rate for individuals aged 15 is 92 percent, 16 is 90 percent and 17 is 80 percent. The drop is faster for low-income individuals. For individuals falling into the bottom two categories of the Colombian poverty index (the SISBEN), the attendance rate for 15 year olds is 84 percent, for 16 year olds is 80 percent and for 17 years olds is 65 percent (Fedesarrollo, 2005). Reflecting these differences, there were 89,000 students who had dropped out of school in 2003. Seventy-four percent of these were classified in the bottom two categories of the SISBEN (Fedesarrollo, 2005).

When surveyed, students claim that the major reason for dropping out is the cost of education. Unlike in many countries, public schools in Colombia are not universally free. Students have to pay to enroll each year and to pay for required items like uniforms, books, and supplies. In fact, 64 percent of dropouts claim that the high cost of education is the main reason for leaving school (Fedesarrollo, 2005). Enrollment fees, uniforms, and school materials make up 90 percent of the costs for low-income individuals, and these monthly costs fluctuates between 24,000 and 50,000 pesos depending on the school and grade (US$ 13 to US$ 22).

III. Research Design

In 2005, the city of Bogota established the Conditional Subsidies for School Attendance ("Subsidios Condicionados a la Asistencia Escolar") program in an effort to improve student retention, lower drop-out rates and reduce child labor. In an effort to improve the program over the basic conditional cash transfer model, the Secretary of Education of the City (Secretaria de Educacion del Distrito, SED) decided to implemented a pilot study in two of the twelve localities in the city. The pilot was to run for a year, and then the results would be used to inform the design of the final program that would operate city-wide.

A. Design of Treatments

Ultimately, three interventions were chosen for the pilot. First, operating as a reference is a basic intervention similar to that used in PROGRESSA/OPPORTUNIDADES. In this basic model, participants would receive 30,000 pesos (approximately US$ 15) as long as the child attended at least 80 percent of the days that month. The payments would be made bi-monthly through a dedicated debit card run by one of the major banks in Colombia. Students would be removed from the program if they failed twice, failed to reach the attendance target in two successive bi-monthly periods, or were expelled from school. Finally, all payments were based on reports provided to the Secretary of Education by the students' principals.

The two additional treatments were experimental variants of this basic intervention aiming to better reach the goals of the program while keeping the cost of each intervention roughly equivalent to the basic intervention.[4] Based on research that suggests that families may face difficulties saving money for students' education (either because of intra-household bargaining, personal discounting issues, or simply high costs of savings), the second treatment (Savings Treatment) varied the timing of the distributions to students' families. Instead of receiving 30,000 pesos a month for reaching the attendance target, students were paid two thirds of this amount on a bi-monthly basis (20,000 pesos or US$10) and the remaining third was held in account. The accumulated funds were then made available to students families during the same period in which students enroll and prepare for the next school year. If students reached the attendance target every month, this treatment would make 100,000 pesos (US$ 50) available to them in December.

Keeping the overall cost of the intervention roughly constant, this treatment differs from the basic intervention with respect to both short-term liquidity constraints and technology to save for longer-term goals. First, because the monthly transfer is reduced, children may attend less often if they face very immediate constraints on school

[4] The amounts, of course, are not the same because the treatments do not account for inflation. Making adjustments to account for inflation probably would have been too complicated to explain to potential registrants.

participation (trading off time spent in school with time spent at work, for example). Second, however, it supplies the accrued funds to families just before they enroll in the next academic year. So, if families' long-term savings constraints are more significant for children's academic participation than the more short-term liquidity constraints, the Savings Treatment could generate both higher attendance and higher re-enrollment rates when compared to the basic treatment.

Rather than manipulate the timing of payments, the third treatment changes the outcome students are being incentivized upon. Instead of providing an incentive to attend school, this treatment provides an incentive to graduate and then to matriculate to a higher education institution. Like the Savings Treatment, this treatment trades off between constraints, but overall the value of the transfer is higher than that of the basic treatment. In the short term, the monthly subsidy is reduced from 30,000 pesos per month to 20,000 pesos. However, upon graduating the students earn the right to receive a transfer of 600,000[5] pesos ($US 300), amounting to 73 percent of the average cost of the first year at a vocational school (823,000 pesos or $US 412). If the student graduates and enrolls in a tertiary institution, they receive the transfer immediately; if they fail to enroll, they can only request the transfer after a year has passed.

Compared to the Basic Treatment, this Tertiary Treatment could reduce attendance rate if students' short-term liquidity constraints are important (because of the lower monthly transfer – as in the Savings Treatment). However, if short-term liquidity constraints are not binding, the Tertiary treatment could stimulate graduation rates and possibly attendance rates (if attendance is viewed as a relevant input into graduation), and could also result in higher levels of matriculation to tertiary institutions.

[5] The amount of 600,000 is equivalent at the yearly savings of the treatment (100,000) time six years between grades 6 and 11. Thus the Tertiary Treatment would be roughly revenue neutral (again, forgetting inflation) if viewed over the full six years that the eventual program would run. Because our study will only evaluate this treatment for students starting in grades 9-11, the total value of this treatment is higher than that of the other treatments.

B. Structure of Randomization

Due to constraints imposed on us by the SED, the assessment of the treatments was divided into two separate experiments located in two very similar localities in Bogota, San Christobal and Suba. Eligible registrants in San Christobal would be randomly assigned between a control group, the Basic Treatment, and the Savings Treatment. Eligible registrants in Suba would be assigned to receive only one of the subsidies, with those who had last completed grades six through eight receiving the Basic Treatment and those who had last completed grades nine through eleven receiving the Tertiary Treatment. This model allows us to directly assess the causal impact of each treatment. It also allows us to directly compare the Savings and Basic Treatments, but it requires us to be careful and ensure the comparability of the localities before comparing the effects of the Tertiary Treatment to the other treatments.

Both experiments were based on an over-subscription model. The city guaranteed enough funds to provide 10,000 with the subsidies, 7,000 in San Christobal and 3,000 in Suba, for three years. To participate, a publicly advertised registration process would be held and if there were more interested children than subsidies, then the subsidies would be allocated to children based on a lottery in each locality.

During January and the beginning of February, the program was advertised in the two localities through posters, newspapers ads, radio spots, loudspeakers in cars, churches, and community leaders, including principals of schools and priests. Potential candidates for the subsidy were registered during 15 days between the end of February and the beginning of March 2005. The registration was conducted in various schools of the two localities. In order to be included in the program, at least one parent / guardian was required to be present at the registration.

In order to be eligible for the program, children had to meet several criteria. First, the potential candidate had to have finished grade 5 and not yet graduated from grade 11. To focus on lower income families, all children's families had to have been classified into the bottom two categories on Colombia's poverty index, the SISBEN.[6] To verify the classification, the student had to present an identification card (which the vast majority of students have). The SISBEN categorization of the household was confirmed online by the SED at the time of registration. In order to eliminate the possibility that families would move to take advantage of the program, only those households that had been classified by the SISBEN system as living in San Christobal or Suba prior to 2004 were eligible to participate in the program.

In all, a total of 17,873 eligible students were registered. Of those, 564 were students who were not currently attending school and were considered for a special version of the subsidy that included remedial assistance and help returning to school.[7] This left 17,309 students eligible for the two experiments: 10, 947 in San Christobal and 6,362 in Suba.

The randomization was publicly conducted on April 4 in each locality. The research team conducted the actual lottery, but in order to ensure transparency of the process, the code was inspected prior the exercise by researchers from the National University. The randomizations were done publicly (projecting the code onto a screen), with representatives of the community, school and local authorities present. The lists of beneficiaries were immediately printed, signed by local officials, and made available to the communities so that parents were able to determine if their children were included.

The randomization was stratified on locality, type of school (public / private), gender, and grade level. Of the 10,000 subsidies, 268 subsidies reserved for the special program for students who had dropped out, and the remaining 9,732 were randomized to the eligible students in our study. Panel A of Table 1 shows the distribution of registrants. In all, 6,875 students from San Cristobal and 2,857 from Suba received one of the treatments. This left 4,072 control students in San Cristobal and 3,505 in Suba, and the students are evenly distributed within grade-gender categories. Finally, while the ratio of assignment is the same within localities, they, of course differ between them. The probability of treatment is San Cristobal was about 63 percent while in Suba the

[6] See Vélez et al (1999) for details for the description of SISBEN. The SISBEN classified households according to 6 levels, 1 being assigned to the poorest. Most of the families in these areas were surveyed in 2003 and 2004.

[7] Unfortunately, this program was never actually implemented, and this portion of the study was discontinued due the lack of an intervention.

probability was 45 percent. We will, of course, have to take account of this difference when pooling samples from both localities.

B. Data

The richness of the available data is one of the major strengths of our study. The data in the current draft of the paper comes from five sources. These include general survey data on all eligible families, data collected specifically for the study, and administrative data collected by the SED.

First, we have the data from the original SISBEN surveys from 2003 and 2004 that contain information on all families eligible to register for the lottery. These surveys were conducted as part of the SISBEN national poverty index – in fact, these are the actual surveys that were used to create the index itself. We have access to all individuals placed into the bottom two SISBEN categories, providing a rich baseline description of the families within the lottery. This provides us with family demographic information, and it also allows us to verify the representativeness of our results by checking that those families who registered for the study were not significantly different from those that did not register. The SISBEN data provide us with several variables at the family level such as schooling level of the household head, physical characteristics of the dwelling, employment status of adults, and family income. It also provides us with individual level variables such as enrollment status at the time of the survey, age, income, and marriage status.[8]

The second source of data comes from the program registration process itself. During this process families had to provide some basic information on the students to ensure eligibility. These data include birth date, gender, last grade completed and year in which that grade was completed. Most of this information was verified through the actual SISBEN data base and when possible, the SED's official records.

---

[8] The obvious challenge of using this data is that families knew that they were being surveyed for the purpose of scoring them on a poverty index. As result, measures of assets and income are probably underestimates of the true values. However, this bias is almost certainly not correlated with the differences investigated in this paper given the timing and purpose of the survey. We use this information for two primary purposes. First we use it to compare registrants to non-registrants, and second we use it as a source of information on the households to which the children in the study belong.

After the randomization, it became clear that students were spread across a large number of schools, but the density was heavily skewed with the majority of students in a smaller number of schools. Based on the available budget, we chose to collect baseline data and the subsequent attendance data in only the 68 schools with the largest number of registered children. This included a total possible sample of 9,768 students. These individuals were chosen from a list of students and the names of the schools that they provided to the SED. Enrollment in these schools was verified by the SED prior to the randomization.

The baseline was conducted between May and July, 2005 and comprised a simple self-administered survey that the students filled out in class. Of the 9,768 students selected for surveying we were able to locate 9,239 students at the time of the baseline survey in the schools that they claimed to attend. The distribution of these students is provided in Panel B of Table 1. Reassuringly, they have a similar distribution to original registrants and again, are equally distributed within grade-gender categories.

Because the baseline was conducted after the randomization, we were unable to use information on any variables that might have changed immediately as a result of the treatments. The baseline instead allows us to narrow down the sample to those children whose provided information was correct and that we could feasibly track down at the end of the study. From the baseline, we use the following: basic demographic variables, a list of friends the students have of the same grade in school, and most importantly, contact information for tracking students during the follow-up survey.

As a fourth source of data, the research team collected during the last quarter of 2005 data on students' attendance through direct observation. For this purpose, the team assembled a group of assistants who randomly visited schools and classes. The assistants directly called the roll of all students and students were marked absent if they were not physically present in the classroom. They visited a total of 1,069 classes in the 68 selected schools for 13 weeks, targeting the same 9,938 students originally chosen for the baseline survey. Because we were able to continue looking for all children selected from the 68 schools, this data set is broader than that used for detailed survey questionnaires as it includes both those students who were found in the baseline and students who, for whatever reason, were not available to be surveyed.

Finally, during February and March of 2006 a follow-up survey was conducted. To ensure that the survey did not preferentially treat students still enrolled in school, we conducted the survey at the household level. For the follow up, the research team located the families of 98.14 percent of the baseline individuals – a total of 8,736 students. The survey is a rich source of information, containing data on the participating students (including academic participation, academic effort; consumption, and labor activities) but also the other children in the household, thereby allowing us to study how the treatments may have affected the allocation of work and resources within households.

C. Analytic Models

We use three basic models to analyze the data. First, we use a simple difference estimator. Second, we also use a difference estimator that includes controls for individual and family characteristics. And finally, we estimate the relationship between attendance and demographic characteristics for control students. We then use this model to estimate what attendance would have been for treatment students without the treatment and for unregistered students had they been observed. In all specifications, we are careful to re-weight the data when pooling results across localities to account for the different treatment assignment ratios.

First, we use a simple difference model to make simple comparisons between different subsets of the sample without controlling for any covariates. These comparisons are intended to assess the comparability of different groups such as the research groups, registrants and non-registrants, etc. When used to compare a given treatment and the respective control group, for example, the specification takes the following form:

$$x_{ij} = \beta_o + \beta_1 Treat_i + \varepsilon_{ij} \qquad (1)$$

To perform this estimate, the data sets containing the treatment group of interest and the respective control group are pooled. The variable x_{ij} represents a particular characteristic of interest for child i in school j. This is regressed on the variable $Treat_i$ which is an indicator variable for whether or not the individual child is in the respective treatment group. The error variable ε_{ij} is indexed with both student and school identifiers because the error terms are allowed to co-vary up to the school level. Finally, the variable β_1 is the estimated difference.

To estimate the effects of the various treatments we use a difference estimator as well, but also include controls for demographic and school characteristics. This model is specified as follows for San Christobal:

$$y_{ij} = \beta_o + \beta_1 Treat1_i + \beta_2 Treat2_i + \delta X_{ijk} + \phi_j + \varepsilon_{ij} \qquad (2)$$

The variables from Equation 1 are defined as before. The variable y_{ij} is the outcome variable of interest. Next, we include two treatment variables that are indicator variables for the specified child receiving the basic and savings treatments, respectively. The coefficients on these indicator variables are the estimates of the effects of the respective treatment. The main difference between this specification and Equation 1 is that this includes as control variables demographic characteristics X_{ijk} at the child and family (k) level as well as fixed effects for each school, ϕ_j. We again allow the error terms to co-vary up to the school level. For Suba, we use a similar equation that contains only one treatment dummy and estimate the model for grades 6-8 and 9-11 separately.

Finally, we use one last specification to estimate what the attendance rates of students who received the treatment would have been without the treatment. Ideally, we would have collected attendance rates of children prior to the randomization. However, we could not have collected this information ourselves because, until the registration process was complete, we had no way of knowing which of the 515,885 eligible students would register. We tried to collect historical attendance rates through the teachers' records, but these records were too often incomplete and when complete, inconsistently kept. To remedy this, we estimate a proxy baseline attendance measure by modeling the control attendance rates using only the available demographic characteristics. Then using the baseline characteristics for treatments students, we their baseline characteristics to project what these students' attendance rates would have been had they not been treated. We then follow a similar procedure for eligible but unregistered students.

To do this, we estimate the following model using only the registered children that did not receive the treatment:

$$y_{ijk} = \beta_o + \delta X_{ijk} + \varepsilon_{ij} \quad (3)$$

The model is estimated using ordinary least squares, and the coefficients and variables are the same as in Equation 2. The only exception, of course, is the omission of the treatment dummies. This equation highlights the fact that this proxy measure is only a linear combination of demographic variables. As such, one interpretation of this variable is as a sufficient statistic for these variables when discussing attendance rates.

IV. Results

We proceed as follows. First, we use the available data from the SISBEN survey to compare the individuals that registered for the program to those who did not and to check comparability between the two localities. Second, for those individuals found at baseline, we compare the students assigned to each research group to ensure that the research groups are balanced at baseline. To make sure that the groups did not become unbalanced due to attrition, we then compare the distribution of students who failed to provide a follow-up survey in each research group. Once we have verified that the groups are indeed still balanced, we then estimate the results of the treatments on the various outcome variables.

A. External Validity

One of the major complaints of randomized evaluations is that, because they often focus on individuals in particular institutions, it remains unclear whether the results can be extrapolated to other populations. In our case, this is a particular concern given that students self-select into the program. However, through the SISBEN surveys, we have access to information on all eligible students living in the two localities, and we can directly compare students whose families registered them for the program to those that did not. The main implication of this comparison is that our results should be applicable to those targeted by the program: poor children currently attending school.

This comparison is presented in Table 2. Each row contains estimates for the indicated demographic variable. Columns 1 and 3 provide the average value for all registered children, and columns 2 and 4 provide the simple difference between registrants and non-registrants using Equation 1. While the size of the sample (515,885 children) is sufficiently large that most differences are statistically significant, they are all very small in magnitude except for those concerning school participation. Families have similar numbers of assets, similar household characteristics, and similar scores on the poverty indexes. Figure 1 shows the entire distribution for our income estimate and similar to the mean, the entire distributions of registrants and non-registrants are comparable.

The main difference is school participation. On average, those registered for the program were more likely to have been attending school when the study was administered (19 and 17 percentage points). There are two reasons for this. First, this particular program targeted students who were already attending school. Second, a primary means of disseminating information about the program was through school principals. This is also born out in Figure 2 where we compare the families using our proxy attendance estimate. Registrants are significantly less likely to be children with similar characteristics to low attending children and much more likely to be similar to those with attendance rates close to 80 percent.

The primary implication of this result is that these results are most applicable to the students for which the interventions were targeted through the eligibility requirements: students who are currently enrolled in school and who have completed at least the fifth grade. However, it also suggests that the program is most attractive to those children with attendance rates close to the target level of 80 percent. Administratively this is attractive because it suggests that the registration process may be a good general strategy for targeting the children most likely to benefit from the program.

Finally, because students are eligible for the Tertiary Treatment only in Suba, we need to make sure that the students in Suba and similar to those in San Cristobal in order to compare properly the magnitudes of the treatment effects. This is done in columns 5 and 6. Column 5 provides a comparison of all eligible children and column 6 provides a comparison of just those children who registered for the lottery. In all cases, these children are very similar, making it reasonable to perform comparisons across localities.

B. Comparison at Baseline

Given that the students who registered for the lottery are representative of all eligible children in the communities, we turn to checking whether or not the randomization succeeded in creating comparable treatment and control groups. This initial comparability is essential for us to be able to attribute future differences between the research groups to the respective treatments.

One problem with the lottery in that not everyone who registered for the program was reachable, most likely because they provided incorrect information at the time of registration. To correct for this and to help us identify the existing sample, we conducted the baseline survey in the 68 schools with the largest number of registrants.

These comparisons are presented in Table 3. As in Table 2, each row displays the comparisons for the indicated demographic variable. Columns 1-4 compare students in San Cristobal and columns 5-8 compare students in Suba. In both localities, the differences are negligible. For San Christobal, columns 2-4 display the simple differences (using Equation 1) between the Basic Treatment and the Control Group, the differences between the Savings Treatment and the Control Group, and finally, the difference between the two treatments, respectively. Almost all of the differences are statistically insignificant and those that are (such as the fact that the Basic Treatment has 3 percent more girls in the sample) statistically significant are economically small.

The same is true for Suba. Columns 5 and 7 respectively show the average control group characteristics for the younger (grades 6-8) and older (grades 9-11) children, respectively. The younger children received the basic treatment, and those selected for the basic treatment are very similar to those in the control group (column 6). Similarly, the older children who received the Tertiary Treatment are similar to the older students who constitute the control group (column 8).

To check for differences in the distribution of children rather than just the mean, we also plotted the distributions. Two of these are shown in Figures 3 and 4. Figure 3 contains a plot of the distribution of household income in the treatment and control groups while Figure 4 contains a plot of our proxy baseline attendance measure. Both figures tell the same story – the distributions are identical.

C. Attrition from Baseline

Comparability at baseline is critical, but once that comparability is established, it is possible that the treatments might cause different types of students to drop out of the study, making the groups incomparable at follow up. We perform two exercises. First, we check the overall attrition rates in each group. If these are sufficiently low, then compositions of the groups cannot significantly change from baseline to treatment even if significantly different types of students attrit. Second, to asses how different the attriters are, we compare the kinds of students attriting in each group using the baseline characteristics of all of the students.

The first two rows of Table 4 provide the exact number of attritors and their percentage in the research group. Column 1 shows the values for the control group and columns 2-4 show the difference from this value and between the two treatment groups for San Cristobal. Columns 5-8 do the same for Suba. Overall, the attrition rate is very low at just less than 2 percent, and the differences in the number of children who dropped out are mostly in the single digits. Given this extremely low rate of attrition, only very large differences could generate changes in the comparability of the research groups.

Panels B through E then estimate these relative comparisons of background characteristics. The control columns (columns 1, 5, and 7) show the difference in characteristics between those students that attrit and those that remain in the sample at follow-up. The difference columns (columns 2-4, 6, and 8) then display the results of a slight modification of Equation 2 to show the difference between the research groups of the relative differences between attritors and stayers.

Again, these differences are relatively minor. The vast majority of the differences are extremely small – for example, the differences in the families as measured through the poverty measures are negligible both in economic and statistical terms. The largest differences occur in the age of the head of the household for San Christobal (3.8 to 6.4

years difference), the age of children in San Cristobal (2.4 years), and the years of education of students in Suba grades 9-11 (1.24 years). Overall the distributions are very similar, and especially given the underlying low rates of attrition, the few differences that do exist are arguably too small to generate confounding changes in the measured outcomes.

D. Results

1. Academic Participation

The fact that the research groups are ultimately comparable allows us to causally attribute any changes in the groups at follow-up to the individual treatments. This allows us to assess families' responses to the various programs by comparing directly the students' who receive the treatments to the control group and to compare directly the different treatment groups. The overall average effects of the treatments combined was to increase verified attendance at school by 2.9 percentage points.

First, we can view the overall effects graphically. The pooled effects of the treatment are depicted in Figure 5 which contains a plot of a kernel density estimate of verified attendance for the treatment and control groups. Based on this graph, the treatment effect seems to operate by reducing the number of students who attend none[9] of the time or between 40 and 70 percent of the time and increases the number of students who attend over 80 percent of the time.

Another way to look at the data is to plot actual attendance rates for each group verse our proxy baseline attendance rates. Using a kernel weighted local polynomial estimator, we plot the relationship of actual measured attendance (on the vertical axis) against the proxy attendance measure (on the horizontal axis). Two results are clear from this graph. First, the treatments have little effect on students with characteristics similar to students who attend over 80 percent of the time without a treatment. Those with characteristics similar to students attending less than 80 percent of the time without the

[9] It is important to note that students with a verified attendance rate of zero may have actually attended school at some point, but just not frequently enough to be caught during one of the visits (up to 13) conducted during the 2007 academic year.

treatment, however, respond significantly, increasing their attendance rates by as much as 10 percent or more. Second, consistent with Figure 5, the effect seems to occur for a wide range of students, not just those who attend slightly less than 80 percent. Those with a proxy attendance rate of 70 percent or more seem to reach the attendance target on average while those attending between 60 and 70 percent attend more despite not reaching the 80 percent attendance target on average (though, of course, they may reach the target at some times).

Dividing up these effects to test for individual effects, we turn to Tables 5 and 6. Table 5 is divided by outcome variables. Panel A contains the results for our most complete outcome measure – verified attendance, and Panel B contains the verified attendance rates just for those students who were found in the follow-up survey. Panel C and D contain self reported attendance and enrollment rates. And finally, Panel E contains the variables pertaining to students in Grade 11. We look at these students individually because in 2006 they would have graduated, and as a result, the outcome variables of interest for these students are unique. In each panel, except for Panel E, the first two rows provide the results for students in grades 6-8 and 9-11 while the third and fourth rows provide estimates for students whose predicted baseline attendance is above and below 80 percent (using Equation 3). Finally, columns 1-3 provide the results for the first experiment in San Cristobal with column 1 providing the average for control students and columns 2 and 3 providing the results for the Basic and Savings Treatments. columns 4-7 provide the results for Suba. Columns 4 and 6 contain the results for the Basic (grades 6-8) and Tertiary (grades 9-11) Treatments while columns 5 and 7 contain the respective controls. All estimates are made using Equation 2.

Turning to Panel A, the individual treatments did cause changes in the verified attendance rates. The first row contains students in grades 6-8 and the second row contains grades 9-11. For grades 6-8 in San Chritobal, the Savings Treatment increases attendance by the same amount as the Basic Treatment (3.7 and 3.8 percentage points respectively), despite the lower monthly transfer. Interestingly, the Basic Transfer has no effect on attendance in Suba. Because only the Basic Treatment was evaluated in Suba, comparing the pooled results to the effects of the Savings Treatment requires us to rely on the comparability of Suba and San Cristobal rather than the experimental design.

However, when we do this, the treatment effect for the Basic Treatment is a statistically insignificant 1.9 percentage points. But although the point estimate is lower than the estimated effect of the Savings Treatment, the difference is not statistically significant. For grades 9-11, the results are different. The results for the Basic and Savings Treatments in San Cristobal are the same, but the results for the Suba experiment (the Tertiary Treatment) are an increase in attendance by 6.1 percentage points, a difference that, when we pool the samples, is statistically different from that of the Basic Treatment.

The next two rows divide students based on the predicted attendance measure estimated from Equation 3. (In other words, we divide the sample based on whether students' baseline characteristics are similar to those characteristics of control group students who either met the attendance target or did not.) As one would expect, the treatment was most effective for students whose projected baseline attendance was below the attendance target. Only the Basic Treatment in San Christobal has a statistically significant effect on students who would have met the target absent the treatment. Students who would not have met the target responded more strongly to each of the treatments.

In Panel B, we focus on just those individuals who were found at school in the baseline survey. From rows one and two, it is clear that primary drivers of the effects reported in Panel A are the students who were not found in the baseline survey. Individuals found in the baseline survey show almost no response to any of the interventions. This is probably due to the fact that, on average, these individuals already attended school more frequently than required by the incentive target.

To check this, we turn to rows 3 and 4 that divide the results by the projected baseline attendance rates. Two interesting results emerge. First, it is clear that those students that found in our baseline survey are different than the average students that we selected from the registration process. Looking at row 4, even conditioning on our baseline attendance measure being less than 80 percent, the students who baseline characteristics fit this classification attended an average of 86 percent of the time. So, conditional on observable characteristics, those students in our baseline and follow-up surveys attend school much more often.

Second, the interventions do affect those students who we would expect to attend less often. In San Christobal, those projected to attend less than 80 percent of the time show an increase in attendance of 1.3 percent due to the Basic Treatment and of 2.1 percent for the Savings Treatment. In Suba, again, there was no response to the Basic Treatment, but the Tertiary Treatment increased attendance by 2.1 percent.

Using this sample, we can also compare our verified attendance rates to the self reported attendance rates that we collected in 2006 using Panel C. There are two significant differences. First, students significantly overestimate their attendance – by about 10 percentage points based on the control averages. Second, the results are inconsistent with those estimated with the verified attendance measures. The only estimates that are close are the estimate effect for the Tertiary Treatment on students who projected baseline attendance would be less than 80 percent. Otherwise, the self-reported estimates show treatment effects for the Basic Treatment in Suba and, strangely, no effects for the San Christobal experiment among students projected to be low attending. Conversely, the self-reported data does show results for the Basic Treatment among students likely to attend schools. It is important to consider that these estimates are from two different years, but it is unlikely that this could account for the different levels of attendance and patterns of estimated effects.

Finally, because students in grade 11 in 2005 should graduate, we divide Panel D and E into grades 6-10 and 11 respectively to take into account the different outcomes variables for these two groups. Panel D provides students self-reported enrollment rates.[10] Only the Savings Treatment and Tertiary Treatment seem to generate significant changes in these rates, but the results for the Savings Treatment are inconsistent. Looking at the Savings Treatment, Enrollment in 2006 is 0.9 percentage points higher for students in grades 6-8 in 2005, but enrollment is 1.4 percentage points less for students in grades 9-11 in 2005. These results are also very sensitive to the specification of the control function. The Tertiary Treatment, however has a more robust 1 percentage point effect for attendance among students with low projected baseline attendance.

[10] We obviously have concerns about the accuracy of self-reported enrollment rates, given the results for the attendance rates.

Panel E contains the results for students who were in the 11th grade in 2005 and should have graduated. And as the control estimates for the first row show – most students (88 and 90 percent) do in fact report graduating. None of the treatments seem to change this rate. However, two of the treatments do have an effect on rates of matriculation to schools of higher education (mostly vocational schools). But the Tertiary Treatments effect of 48 percentage points is dramatically higher than the effect of a statistically insignificant 4.9 percentage points for the Basic Treatment. The Savings Treatment also increased the enrollment rate by more than the Basic Treatment (9.5 percentage points), but the difference between the two treatment effects is not statistically significant.

2. Other Outcomes

While academic participation is the main outcome of interest, we also collected other outcome variables which are presented in Table 6. Because the results were relatively similar across grades, we pooled all of the grades in Panels A and B to investigate academic effort and consumption. The columns are defined as in Table 5 and all estimates are again done using Equation 2.

In general, the treatments have a mixed effect on our measures of academic behavior. The only significant effect, both statistically and practically, is that the Tertiary Treatment increases the time spent on homework by a half an hour a week. The other two treatments do not have any affect. Given the magnitude of the changes in attendance, it would be surprising to see changes in grades as a result of only increased exposure to school, but children may have exerted more effort when they did go to school as a result of the treatment. The results in rows 2 and 3 suggest otherwise. In neither self-reported or verified grades do we see change in grades induced by the treatments. However, the Savings Treatments does seem to increase the number of students who matriculate to the next grade by 2 percentage points.

The effects on food consumption are small but largely uniform across the treatments. As shown in Panel B, all the treatments increase the number of meals children have eaten over the last three days by about 0.15 to 0.24 of a meal. Similarly, the number of meals with a common source of protein increases by 0.16 to 0.18 meals over 3 days. Finally, we find no effect of any of the treatments on textbook ownership.

Panels C and D contain the results for hours of work both in the labor market and at home. Panel C contains the results for students in grades 6 through 8. In both the Savings and the Basic treatment (in both Suba and San Christobal), there is very little change in the number of hours worked. This is true for both hours worked for pay and hours worked without pay.

Panel D contains the results for older children. As rows 1-3 of Panel D indicate, the Basic and Savings Treatments still have little effect on the hours that students work. However, the Tertiary Treatment seems to have a large effect, reducing total hours by 3 hours a week and hours for pay by over 2.25 hours a week.

This treatment however also affected other children in the family. The remaining rows in Panel D show the results for other children of school age (between 6 and 21 years of age)[11] in the family. While the basic and savings treatment still have no effect, there seems to be a reallocation of work within the family caused by the Tertiary Treatment. The primary result of which is to increase the total work effort for pay by other children together by about 1.68 hours a week. Dividing this between children currently attending and those not attending school, it seems that both groups bear some of the additional burden.

3. Heterogeneity in the Treatment Effects

Another important dimension of these incentive programs is their relative impact on different types of students. We study possible heterogeneity by estimating Equation 2 for different subsets of the sample and estimated the differences for three main outcomes: verified attendance, number of meals, and the number of textbooks. However, in order to maintain sufficient sample size within subsets of the data, we pooled the data and estimate the average effects across all three treatments.

[11] This age group is arbitrary. We find similar results using other upper bounds such as 15 or 18 years of age.

These results are in Table 7. In Panel A, the columns are presented in groups of two. Columns 1 and 2 provide the results for verified attendance. Columns 3 and 4 display results for the number of meals and columns 5 and 6 show the results for the number of textbooks. The first column shows the average value for the control group and the second contains the difference between the combined treatments and control groups. The first row contains the overall weighted average values for all students in the sample.

Rows 2 and 3 contain the results divided by our proxy baseline attendance measure. The sample is divided into two groups with those whose projected attendance without the program would be more than 80 percent on the second row and those below 80 percent on the second row. As we already saw before, only students whose attendance without the program would be under 80 percent responded to the program in behaviors related to school. Those who would have attended under 80 percent increase their attendance by 3.8 percentage points while those who would have already been attending enough to meet the threshold increase their attendance by a statistically insignificant 0.1 percentage points. Interestingly, families with initially low attendance are also the ones most likely to record an increase the number of textbooks owned by the child (5.8 percentage points vs. 1.7 percentage points for those with attendance over 80 percent). Not surprisingly, the cash transfers do seem to more generally increase the consumption of non-academically related items like food. This is shown in columns 3 and 4. Unlike the academically related outcomes, all families increase the number of meals the children received by a similar amount (0.238 meals for those attending over 80 percent vs. 0.145 for those not initially meeting the target). This does demonstrate that outside of the incentive effects, the small transfers from the program can have a measurable effect on general wellbeing.

The next two rows divide the sample by gender. There seems to be a large difference in the responses of boys' and girls' attendance patterns (4.3 percentage points to 1.6 percentage points).[12] Finally, there seems to be no difference in the number of meals or textbooks received by boys or girls.

[12] Breaking the sample down further, it seems that this result is primarily driven by a large difference in the responses between boys and girls between 6 and 8th grade for the Savings Treatment.

Next, we break down the sample into terciles using our measure of family income. The families at the top of the income distribution show the strongest response (4.9 percentage points on the verified attendance measure) while those at the bottom show no measurable response at all (1.4 percentage points of verified attendance). This suggests that the transfer may just be too small to make a difference in the lives of the poorest families. All of these students are poor, but apparently, those who are less poor have the capacity to take advantage of this program while those with fewer resources do not.

Next, we divide the sample by the number of years the head of the household had been in school. Along this dimension, no clear pattern emerges. Students in each tercile experience similar changes in verified attendance rates and changes in the number of meals. The change in the number of textbooks, however, seems to be strongest for families with the most educated household heads. Perhaps more educated household heads are more likely to see the value of spending additional money on textbooks.

Finally, Panel B breaks down the tertiary treatment using three outcome variables displayed in columns 1 and 2, 3 and 4, and 5 and 6 respectively: hours worked by registered student for pay, hours worked by the registered student without pay, and finally the hours worked by other school aged children for pay. In this case we divide the sample by family income and household education.

Three main findings emerge. First, the decrease in hours worked by the registered student is almost always matched by a concomitant increase in the number of hours worked by other children in the household. Second, when the sample is divided by family income, the pattern of the poorest families responding least is again apparent, both in the number of hours worked by registered students and in the increased hours worked by other children. Finally, while this differences is not statistically significant, it seems that the reduction in hours worked by the registered student is strongest in the least educated households; the same is true (and the pattern is clearer) for the increase in hours worked by other children in the household.

V. Conclusion

This project demonstrates that experiments involving conditional cash transfer programs can be used to understand how variation in structure these programs may affect educational and related outcomes for targeted children and their family Among other things, our results suggest that, in this environment at least, short-term liquidity issues are not significant enough that decreasing the monthly subsidy by a third has any effect on attendance. In addition, the results suggest that incentives focused on completing academic goals (such as graduation) may have larger effect on attendance, than those that only condition on attendance.

The Savings and Tertiary Treatments improved enrollment in higher education, but the Tertiary Treatment proved much more effective (a 48 percent increase compared to a 9.5 percent increase). The Tertiary Treatment also increased enrollment among low attending students, and increase the amount of time spent on homework by a half and hour a week on average.

Exploiting the within family variation in our data suggests that families also respond to the relative opportunities available to their children and reallocated labor market responsibilities to allow recipients to take advantage of the Tertiary Treatment. Specifically, while recipients worked fewer hours on average (3 per week), their school aged siblings covered this reduction by working on average 1.8 more hours a week.

Finally, looking within subsets of our sample there is also very strong variation in the treatment effects across groups. Not surprisingly, the treatments we experimented with provided stronger incentives for children who are most likely to have attendance rates lower than the subsidies' target. Girls also seem to respond less to boys, though the exact reason is unclear. Finally, the effects of the treatments along all dimensions seem to be concentrated among families in the top two terciles of our sample. The specific reasons as to why the poorest of the poor may not be able to take advantage of these programs are an important topic for future work.

References

Angrist, Joshua and Alan Krueger (1991) “Does Compulsory School Attendance Affect Schooling and Earnings?” *Quarterly Journal of Economics*. 106(4): 979-1014.

Attanasio, Orazio, Erich Battistin, Emla Fitzsimons, Alice Mesnard and Marcos Vera-Hernández (2005). “How Effective are Conditional Cash Transfers?: Evidence from Colombia.” The Institute of Fiscal Studies Briefing Note No. 54.

Attanasio, Orazio, Emla Fitzsimmons and Ana Gomez (2005). “The Impact of a Conditional Education Subsidy on School Enrollment in Colombia.” The Institute of Fiscal Studies, Report Summary Familias 01.

Behrman, Jere R., Pilali Sengupta and Petra Todd (2005). "Progressing Through PROGRESA: An Impact Assessment of a School Subsidy Experiment in Mexico." *Economic Development and Cultural Change.* 54(1): 237-275.

Barrera-Osario, Felipe, Leigh L. Linden, Miguel Urquiola (2007) “The Effects of User Fee Reductions on Enrollment: Evidence form a Quasi-Experiment,” Columbia University Department of Economics Mimeo.

Cardoso, Eliana and André Portela Souza (2004). “The Impact of Cash Transfers on Child Labor and School Attendance in Brazil.” Vanderbilt University Working Paper No. 04-W07.

Chaudhury, Nazmul and Dilip Parajuli (2006). “Conditional Cash Transfer and Female Schooling: Impact of the Female School Stipend Program on Public School Enrollments in Punjab, Pakistan.” Work Bank Policy Research Working Paper 4102.

Duflo, Esther (2001) “Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from and Unusual Policy Experiment,” *American Economic Review*. 91(4): 795-813.

Fedesarrollo (2005) “Projecto Piloto, Subsidios Condicionados a Asistencia Escolar en Bogota: Diseño del Piloto y la Evaluación de Impacto (Informe Final)” *Mimeo*, Fedesarrollo, Bogota, Colombia.

Filmer, Deon and Norbert Schady (2006). “Getting Girls into School: Evidence from a Scholarship Program in Cambodia.” World Bank Policy Research Working Paper 3910.

Glewwe, Paul and Pedro Olinto (2006). “Evaluating the Impact of Conditional Cash Transfers on Schooling: An Experimental Analysis of Honduras PRAF Program. Final Report for USAID.” International Food Policy Research Institute.

Krueger, A. and M. Lindahl (2001). “Education for Growth: Why and For Whom?” *Journal of Economic Literature*, Vol. 39, No. 4, 1101-1136.

Levy, Dan and Jim Ohls (2006). “Evaluation of Jamaica’s Path Program: Final Report.” Mathematica.

Maluccio, John A. and Rafael Flores (2005). Impact Evaluation of a Conditional Cash Transfer Program: The Nicaraguan Red de Protección Social.” Washington, D.C.: International Food Policy Research Institute.

Muralidharan, Kartik and Venkatesh Sundararaman (2006) “Teacher Incentives in Developing Countries: Experimental Evidence from India,” *Working Paper.* Harvard University Department of Economics.

Pitt, Mark, Shahidur Khandker and Nubuhiko Fuwa (2003). “Subsidy to Promote Girls’ Education: The Female Stipend Program in Bangladesh.” Mimeo.

Schady, Norbert and Maria Caridad Araujo (2006). “Cash Transfers, Conditions, School Enrollment, and Child Work: Evidence from a Randomized Experiment in Ecuador.” World Bank Policy Research Working Paper 3930.

Schultz, T. Paul (2004) “School Subsidies for the Poor: Evaluating the Mexican Progresa Poverty Program”, *Journal of Development Economics*, 74(1):199-250

United Nations (2006) *Deevelopment Goals Report 2006,* United Nations, New York.

Vermeersch, Christel and Michael Kremer, “School Meals, Educational Achievement, and School Competition: Evidence from a Randomized Evaluation” World Bank Policy Re-search Working Paper: No. 3523, 2005.

Villatoro, Pablo (2005). “Conditional Cash Transfer Programmes: Experiences from Latin America.” *CEPAL Review* 86: 83-96.

World Bank (2006). *World Development Indicators 200,* The World Bank, Washington.

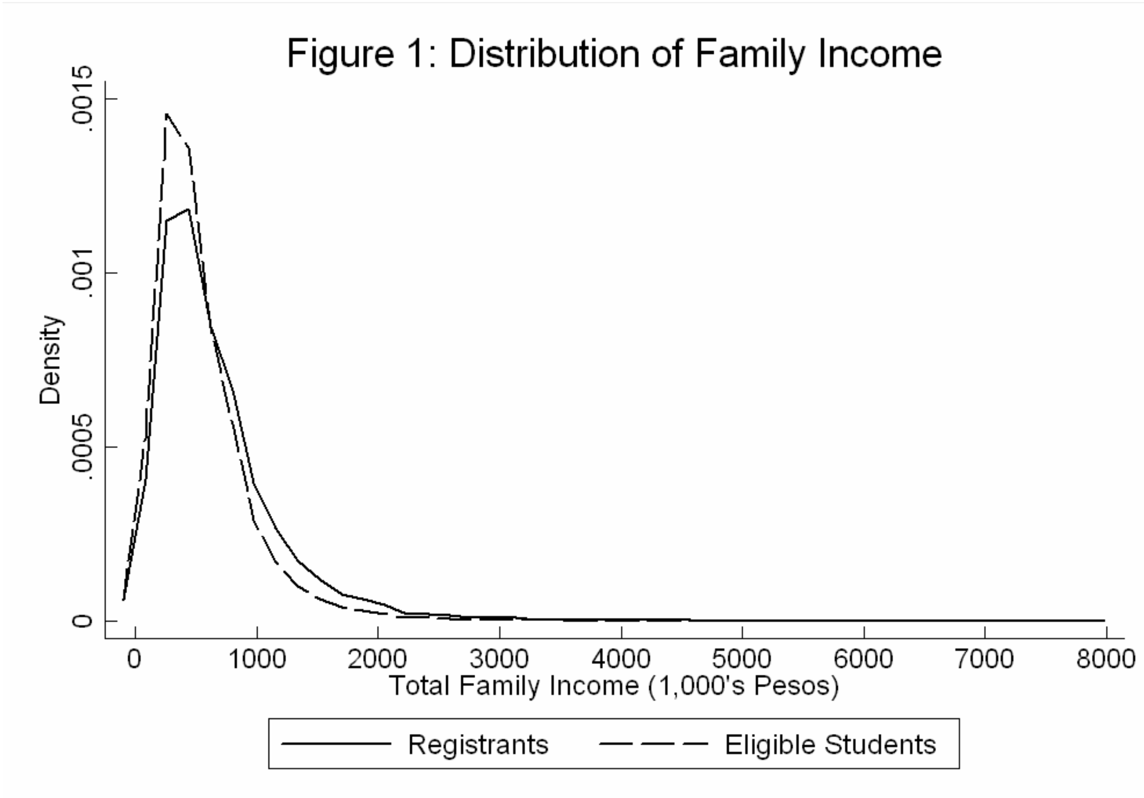

Figure 1: Distribution of Family Income
Density
0
.0005
.001
.0015
0
1000
2000
3000
4000
5000
6000
7000
8000
Total Family Income (1,000's Pesos)
Registrants
Eligible Students

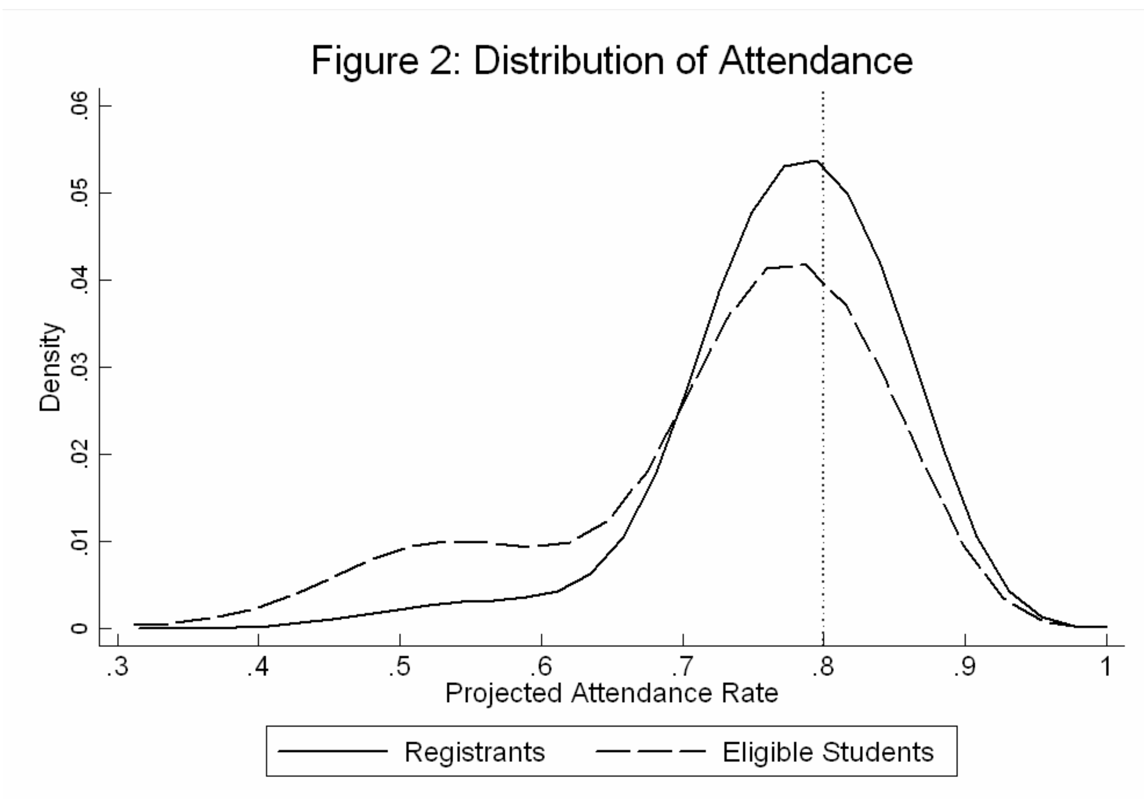

Figure 2: Distribution of Attendance
Density
0
.01
.02
.03
.04
.05
.06
.3
.4
.5
.6
.7
.8
.9
1
Projected Attendance Rate
Registrants
Eligible Students

Figure 3: Distribution by Family Income at Baseline

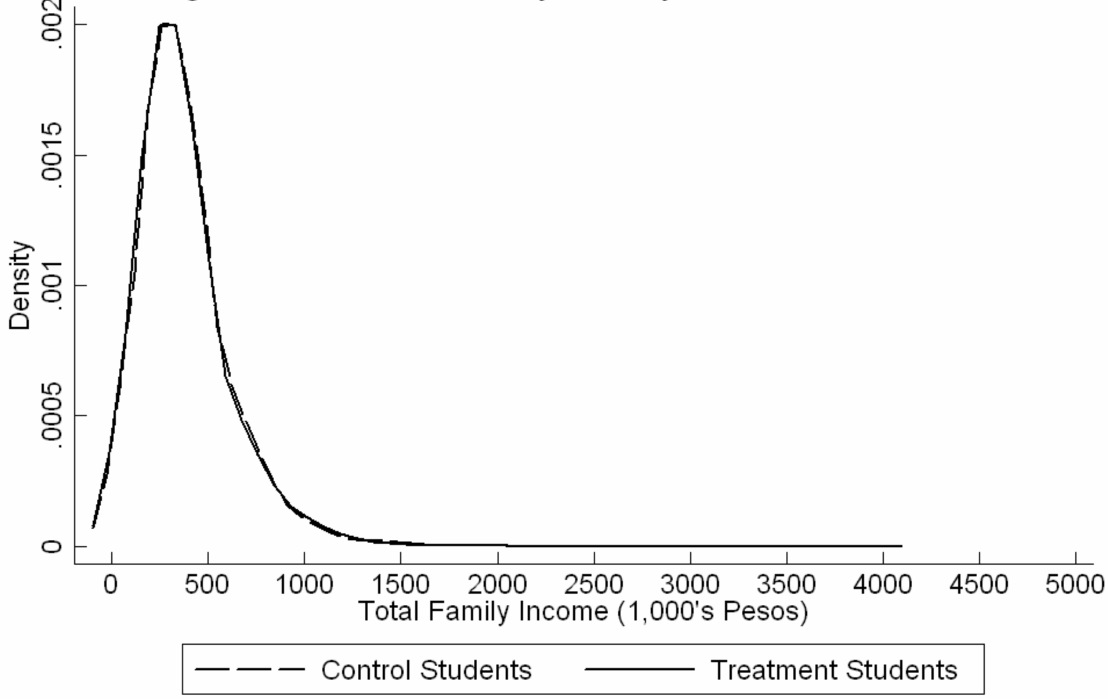

Figure 4: Distribution of Attendance at Baseline

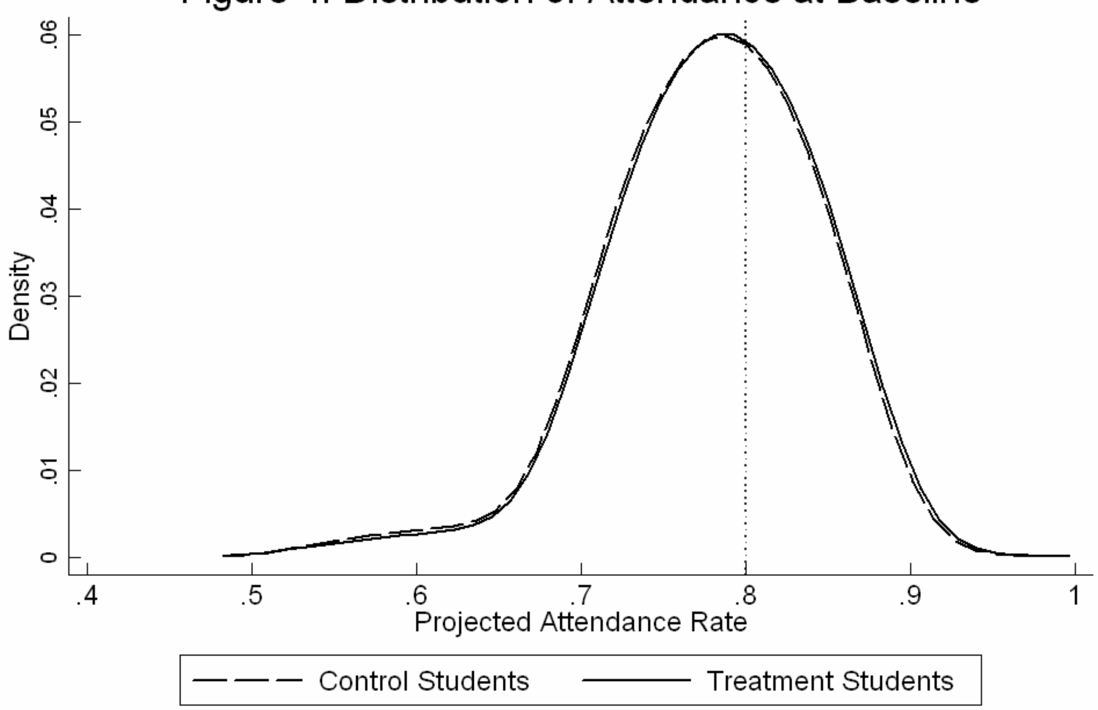

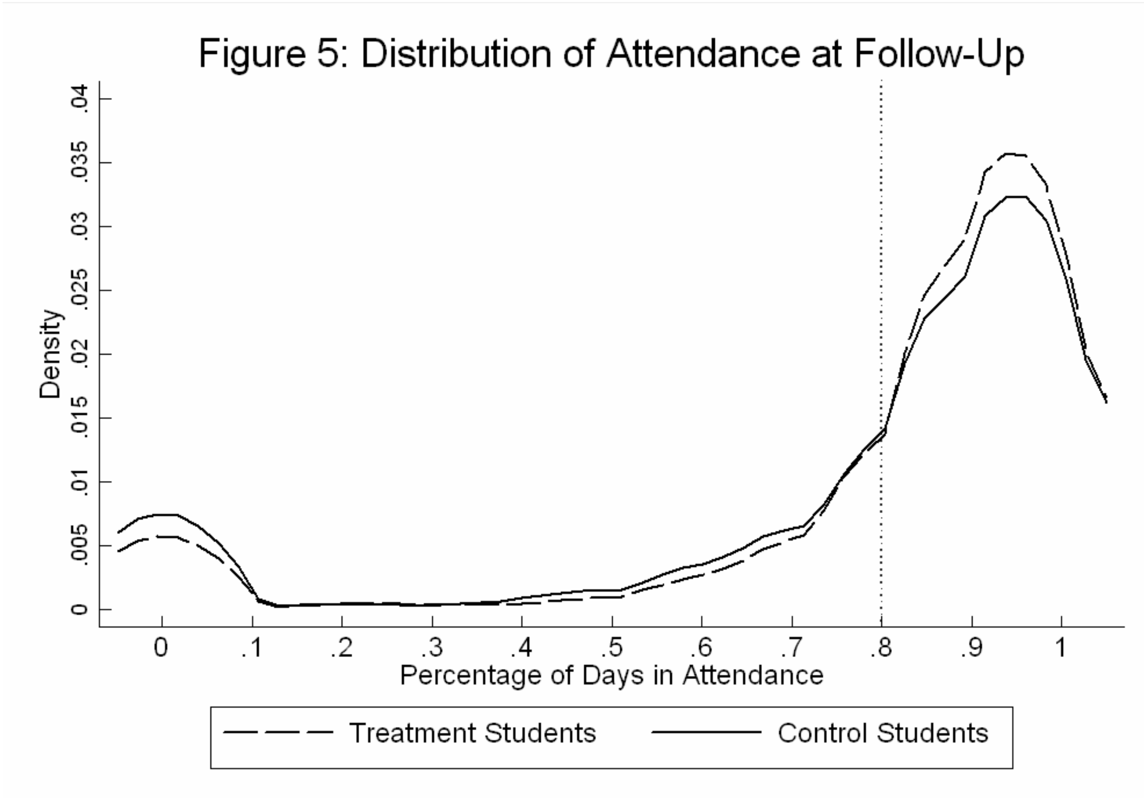
Figure 5: Distribution of Attendance at Follow-Up
Density
0
.005
.01
.015
.02
.025
.03
.035
.04
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1
Percentage of Days in Attendance
Treatment Students
Control Students

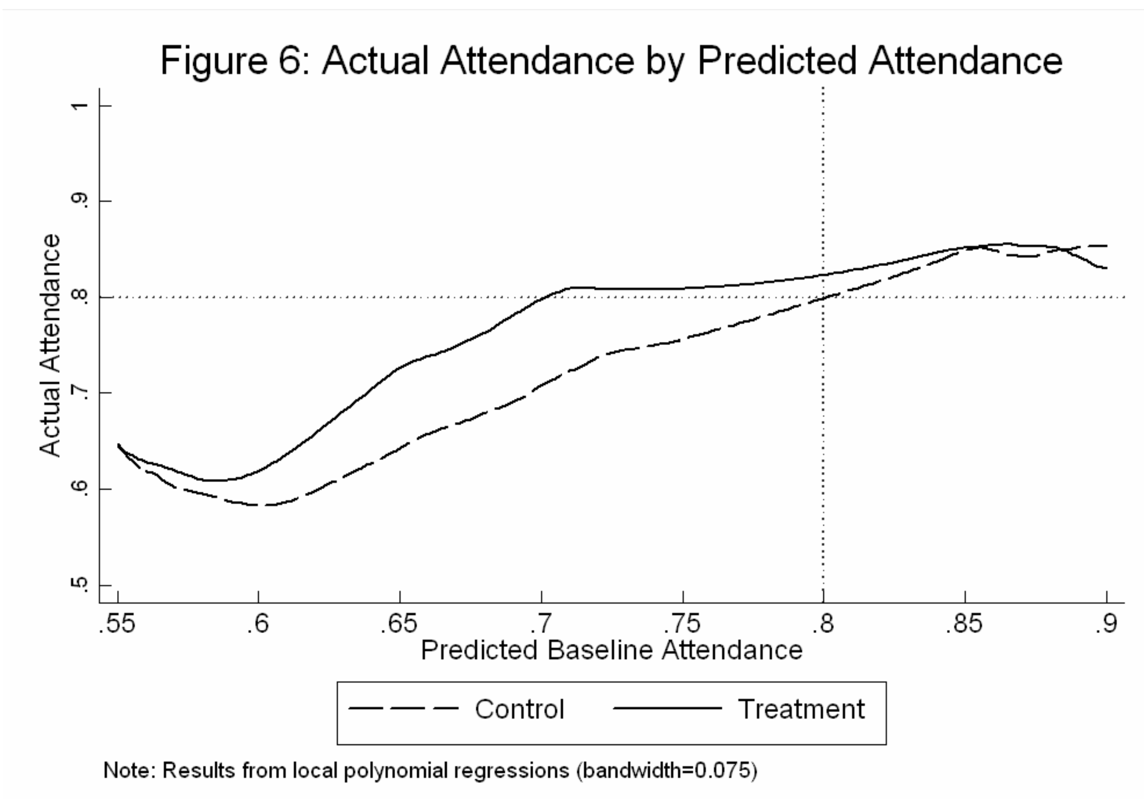
Figure 6: Actual Attendance by Predicted Attendance
Actual Attendance
.5
.6
.7
.8
.9
1
.55
.6
.65
.7
.75
.8
.85
.9
Predicted Baseline Attendance
Control
Treatment
Note: Results from local polynomial regressions (bandwidth=0.075)

Table 1: Distribution of Subjects by Research Groups

		San Cristobal			Suba		
Grade		**Control**	**Treat 1**	**Treat 2**	**Control**	**Treat**	**Total**
Panel A: All Registrants							
Six	**Male**	430	362	362	374	306	**1834**
	Female	475	402	402	386	315	**1980**
Seven	**Male**	402	341	341	350	286	**1720**
	Female	401	340	340	335	274	**1690**
Eight	**Male**	380	319	319	346	283	**1647**
	Female	360	305	305	310	253	**1533**
Nine	**Male**	342	289	290	300	244	**1465**
	Female	329	277	277	262	214	**1359**
Ten	**Male**	310	261	261	269	218	**1319**
	Female	266	226	226	223	181	**1122**
Eleven	**Male**	213	178	178	212	172	**953**
	Female	164	137	137	138	111	**687**
Total		**4072**	**3437**	**3438**	**3505**	**2857**	**17309**
Panel B: Students Found at Baseline							
Six	**Male**	211	201	200	195	178	**985**
	Female	201	237	208	187	170	**1003**
Seven	**Male**	218	217	207	188	174	**1004**
	Female	199	204	196	160	154	**913**
Eight	**Male**	183	171	174	188	164	**880**
	Female	179	187	186	148	141	**841**
Nine	**Male**	176	174	164	142	121	**777**
	Female	155	174	147	100	101	**677**
Ten	**Male**	167	152	152	124	122	**717**
	Female	147	135	125	89	86	**582**
Eleven	**Male**	95	109	108	99	95	**506**
	Female	87	80	70	63	54	**354**
Total		**2018**	**2041**	**1937**	**1683**	**1560**	**9239**

Table 2: Comparison of Registered and Eligible Students

Demographic Variable	San Cristobal Eligible Children	San Cristobal Registrants - Eligible	Suba Eligible Children	Suba Registrants - Eligible	San Cristobal - Suba Eligible Children	San Cristobal - Suba Registered Children
Panel A: Indexes of Household Assets						
Posessions	1.91	0.02	1.83	0.02	0.08***	0.08***
	(1.09)	(0.01)	(1.00)	(0.01)	(0.01)	(0.02)
Utilities	4.77	-0.10***	4.85	-0.10***	-0.08***	-0.08***
	(1.40)	(0.01)	(1.35)	(0.02)	(0.01)	(0.02)
Durable Goods	1.5	-0.14***	1.67	-0.10***	-0.18***	-0.20***
	(0.94)	(0.01)	(0.92)	(0.01)	(0.01)	(0.01)
Physical Infrastructure	11.9	-0.25***	12.23	-0.25***	-0.34***	-0.32***
	(1.79)	(0.02)	(1.58)	(0.02)	(0.01)	(0.03)
Panel B: Individual Characteristics						
Age	15.16	-1.47***	15.08	-1.38***	-0.02	-0.01
	(3.33)	(0.03)	(3.33)	(0.04)	(0.02)	(0.03)
Gender	0.5	0.01	0.5	-0.01*	0	0.02**
	(0.50)	(0.01)	(0.50)	(0.01)	(0.00)	(0.01)
Married	4.84	0.15***	4.84	0.14***	-0.00***	0.01***
	(0.77)	(0.01)	(0.77)	(0.01)	(0.00)	(0.00)
Attending School	75.61	18.96***	76.15	16.66***	0.84***	1.75***
	(42.95)	(0.43)	(42.62)	(0.54)	(0.24)	(0.38)
Years of Education	6.33	-0.69***	6.5	-0.66***	-0.22***	-0.21***
	(3.08)	(0.03)	(3.08)	(0.04)	(0.02)	(0.03)
Panel C: Household Characteristics						
Single Head	0.33	-0.03***	0.31	-0.03***	0.02***	0.03***
	(0.47)	(0.00)	(0.46)	(0.01)	(0.00)	(0.01)
Age of Head	45.97	0	44.77	0.36***	1.15***	0.83***
	(11.05)	(0.12)	(10.03)	(0.13)	(0.06)	(0.15)
Years of Ed, Head	5.91	-0.35***	6.33	-0.47***	-0.43***	-0.30***
	(3.09)	(0.03)	(3.16)	(0.04)	(0.02)	(0.05)
People in Household	5.27	0.14***	5.07	0.14***	0.22***	0.22***
	(2.06)	(0.02)	(1.89)	(0.02)	(0.01)	(0.03)
Member under 18	2.14	0.44***	2.03	0.48***	0.13***	0.07***
	(1.49)	(0.02)	(1.39)	(0.02)	(0.01)	(0.02)
Panel D: Poverty Measures						
Estrato	1.49	-0.04***	1.61	-0.02**	-0.13***	-0.14***
	(0.82)	(0.01)	(0.78)	(0.01)	(0.00)	(0.01)
SISBEN Score	12.89	-1.18***	14.28	-1.14***	-1.46***	-1.42***
	(5.01)	(0.05)	(4.74)	(0.06)	(0.03)	(0.07)
Household Income	437.04	-70.22***	482.96	-83.32***	-48.23***	-32.81***
(1,000 Pesos)	(282.34)	(2.92)	(295.62)	(3.83)	(1.66)	(3.78)

* significant at 10 percent level, ** at 5 percent level, and *** at 1 percent level
Standard errors are clustered at the school level.

Table 3: Baseline Comparison of Students

Demographic Variable	San Cristobal Control Average	San Cristobal Treat 1 - Control	San Cristobal Treat 2 - Control	San Cristobal Treat 1 - Treat 2	Suba (Grade 6-8) Control Average	Suba (Grade 6-8) Treat - Control	Suba (Grade 9-10) Control Average	Suba (Grade 9-10) Treat - Control
Panel A: Indexes of Household Assets								
Posessions	1.94	0.06*	<0.01	0.06*	1.84	-0.03	1.93	0.03
	(1.11)	(0.03)	(0.03)	(0.03)	(1.00)	(0.05)	(1.00)	(0.06)
Utilities	4.67	-0.03	0.04	-0.06	4.67	0.06	4.86	0.08
	(1.42)	(0.05)	(0.05)	(0.04)	(1.40)	(0.05)	(1.31)	(0.06)
Durable Goods	1.37	-0.03	0.02	-0.05**	1.53	0.02	1.62	0.07
	(0.88)	(0.03)	(0.03)	(0.02)	(0.87)	(0.04)	(0.83)	(0.05)
Physical Infrastructure	11.64	-0.11**	0.01	-0.12**	11.9	0.03	12.12	0.01
	(1.73)	(0.05)	(0.04)	(0.06)	(1.49)	(0.05)	(1.38)	(0.08)
Panel B: Individual Characteristics								
Age	14.36	0.01	-0.22	0.23	12.66	0.07	15.58	0.21
	(5.50)	(0.12)	(0.21)	(0.19)	(3.87)	(0.16)	(4.28)	(0.32)
Gender	0.49	0.02	-0.01	0.03**	0.48	0.01	0.42	<0.01
	(0.50)	(0.01)	(0.01)	(0.01)	(0.50)	(0.02)	(0.49)	(0.04)
Married	4.92	-0.01	0.01	-0.02	4.98	-0.02	4.95	-0.01
	(0.53)	(0.01)	(0.02)	(0.02)	(0.30)	(0.01)	(0.41)	(0.03)
Years of Education	5.6	-0.09**	-0.06	-0.03	4.67	0.06	7.45	-0.07
	(1.88)	(0.04)	(0.06)	(0.05)	(1.26)	(0.06)	(1.24)	(0.07)
Panel C: Household Characteristics								
Single Head	0.29	<0.01	<0.01	<0.01	0.27	<0.01	0.25	0.03
	(0.46)	(0.01)	(0.01)	(0.02)	(0.44)	(0.01)	(0.43)	(0.02)
Age of Head	45.75	-0.13	0.09	-0.22	44.68	-0.2	45.72	0.82**
	(10.32)	(0.27)	(0.36)	(0.30)	(9.10)	(0.42)	(8.52)	(0.33)
Years of Ed, Head	5.58	-0.07	-0.07	0	5.68	0.06	5.8	-0.12
	(2.90)	(0.09)	(0.08)	(0.10)	(2.78)	(0.09)	(2.98)	(0.13)
People in Household	5.39	0.03	0.02	0.01	5.25	0	5.13	0.09
	(1.91)	(0.06)	(0.06)	(0.05)	(1.75)	(0.09)	(1.62)	(0.08)
Member under 18	2.59	0.06	0.03	0.03	2.7	0.01	2.37	0.03
	(1.32)	(0.04)	(0.04)	(0.04)	(1.29)	(0.07)	(1.18)	(0.09)
Panel D: Poverty Measures								
Estrato	1.45	-0.01	<0.01	-0.01	1.57	0.01	1.65	<0.01
	(0.81)	(0.03)	(0.03)	(0.02)	(0.82)	(0.03)	(0.76)	(0.04)
SISBEN Score	11.76	-0.25*	-0.15	-0.09	13.11	-0.01	13.51	0.23
	(4.63)	(0.14)	(0.14)	(0.13)	(4.37)	(0.12)	(4.25)	(0.26)
Household Income	364.46	-2.12	5.07	-7.19	389.74	0.35	396.89	3.42
(1,000 Pesos)	(235.54)	(7.36)	(8.67)	(6.99)	(223.17)	(9.33)	(228.04)	(12.49)

* significant at 10 percent level, ** at 5 percent level, and *** at 1 percent level
Standard errors are clustered at the school level.

Table 4: Attrition from Baseline Survey

Demographic Variable	San Cristobal Control Average	Treat 1 - Control	Treat 2 - Control	Treat 1 - Treat 2	Suba (Grade 6-8) Control Average	Treat - Control	Suba (Grade 9-10) Control Average	Treat - Control
Panel A: Attrition Rate								
Number Attritors	44	-2	-15	13	18	3	8	3
Percentage of Baseline	0.02	0	-0.01*	0.01	0.02	0	0.01	0.01
	(0.15)	(0.01)	(0.00)	(0.01)	(0.13)	(0.00)	(0.11)	(0.01)
Panel B: Indexes of Household Assets								
Posessions	-0.31*	-0.05	-0.21	0.15	-0.49**	0.05	-0.31	0.17
	(0.17)	(0.17)	(0.22)	(0.17)	(0.24)	(0.19)	(0.36)	(0.44)
Utilities	-0.75***	0.46	0.40*	0.06	-0.63*	-0.51	0.14	-1.75***
	(0.22)	(0.30)	(0.24)	(0.29)	(0.34)	(0.40)	(0.47)	(0.50)
Durable Goods	-0.21	-0.01	-0.08	0.06	-0.55***	0.08	-0.50*	-0.01
	(0.13)	(0.22)	(0.20)	(0.20)	(0.21)	(0.24)	(0.30)	(0.28)
Physical Infrastructure	-0.31	-0.2	-0.36	0.16	-0.16	-0.39	0.13	-1.38**
	(0.26)	(0.36)	(0.42)	(0.46)	(0.35)	(0.45)	(0.49)	(0.65)
Panel C: Individual Characteristics								
Age	2.63***	-2.4	-3.02*	0.62	-0.13	0.53	0.47	-0.9
	(0.84)	(2.09)	(1.59)	(1.24)	(0.95)	(0.40)	(1.49)	(0.63)
Gender	0.11	-0.14	-0.12	-0.02	0.23*	-0.29***	-0.30*	0.24
	(0.08)	(0.11)	(0.13)	(0.11)	(0.12)	(0.10)	(0.18)	(0.24)
Married	-0.06	0.03	0.13	-0.1	0.02	0.02	0.04	0.02
	(0.08)	(0.14)	(0.09)	(0.10)	(0.07)	(0.01)	(0.14)	(0.03)
Years of Education	-0.28	0.36	0.17	0.18	0.04	0.6	1.32***	-1.24**
	(0.29)	(0.44)	(0.66)	(0.50)	(0.31)	(0.37)	(0.44)	(0.49)
Panel D: Household Characteristics								
Single Head	-0.09	0.20**	0.19*	0.01	0.09	0.08	0.12	0.05
	(0.07)	(0.09)	(0.11)	(0.14)	(0.11)	(0.14)	(0.15)	(0.16)
Age of Head	-1.45	-3.80**	2.55	-6.35***	-3.63	3.25	0	-4.73
	(1.57)	(1.72)	(2.17)	(2.40)	(2.23)	(2.33)	(3.04)	(3.30)
Years of Ed, Head	0.89**	-0.81	-0.57	-0.23	0.55	-1.63	0.6	-1.2
	(0.44)	(0.63)	(0.67)	(0.81)	(0.68)	(1.00)	(1.06)	(1.36)
People in Household	0.06	-0.44	-0.19	-0.25	-0.26	-0.62	-0.38	-0.4
	(0.29)	(0.35)	(0.36)	(0.32)	(0.43)	(0.50)	(0.58)	(0.80)
Member under 18	-0.09	0.25	0.22	0.03	0.25	-0.47	0.14	-0.36
	(0.20)	(0.38)	(0.38)	(0.24)	(0.32)	(0.37)	(0.42)	(0.71)
Panel E: Poverty Measures								
Estrato	-0.34***	0.17	0.25	-0.09	-0.04	-0.45	0.1	-1.13***
	(0.12)	(0.19)	(0.21)	(0.19)	(0.20)	(0.29)	(0.27)	(0.30)
SISBEN Score	-1.56**	0.9	1.1	-0.21	-1.09	-1.76	-0.82	-4.00**
	(0.71)	(0.98)	(1.01)	(1.24)	(1.07)	(1.28)	(1.51)	(1.67)
Household Income	-0.86	-49.57	-42.21	-7.36	-104.63**	74.41	-87.77	43.52
(1,000 Pesos)	(35.93)	(48.92)	(53.80)	(44.25)	(52.99)	(61.74)	(81.20)	(65.96)

* significant at 10 percent level, ** at 5 percent level, and *** at 1 percent level
Standard errors are clustered at the school level.

Table 5: Academic Participation Outcomes

Outcome Variable	San Christobal			Suba, Grades 6-8		Suba, Grades 9-11	
	Control Average	Basic - Control	Savings - Control	Control Average	Basic - Control	Control Average	Tertiary - Control
Panel A: Verified Attendance							
Grades 6-8	0.776	0.037***	0.038***	0.764	0.003		
	(0.008)	(0.012)	(0.012)	(0.009)	(0.016)		
Grades 9-11	0.79	0.024**	0.040***			0.789	0.061***
	(0.010)	(0.010)	(0.012)			(0.011)	(0.015)
Baseline Att >= 0.8	0.824	0.022*	0.011	0.814	-0.008	0.833	-0.035
	(0.009)	(0.013)	(0.012)	(0.017)	(0.016)	(0.020)	(0.028)
Baseline Att < 0.8	0.759	0.039***	0.053***	0.749	0.007	0.778	0.077***
	(0.008)	(0.010)	(0.011)	(0.011)	(0.020)	(0.013)	(0.018)
Panel B: Verified Attendance if Followup							
Grades 6-8	0.874	0.009	0.01	0.84	-0.002		
	(0.005)	(0.006)	(0.006)	(0.007)	(0.008)		
Grades 9-11	0.869	0.009	0.016			0.857	0.019
	(0.006)	(0.008)	(0.010)			(0.008)	(0.013)
Baseline Att >= 0.8	0.891	0.003	-0.004	0.856	0.007	0.869	0.004
	(0.005)	(0.006)	(0.007)	(0.014)	(0.011)	(0.015)	(0.029)
Baseline Att < 0.8	0.861	0.013*	0.021**	0.835	-0.004	0.853	0.021*
	(0.005)	(0.007)	(0.008)	(0.008)	(0.010)	(0.009)	(0.012)
Panel C: Self Reported Attendance							
Grades 6-8	0.954	0.010**	0.009**	0.961	0.010**		
	(0.004)	(0.005)	(0.005)	(0.003)	(0.004)		
Grades 9-11	0.963	-0.004	0.003			0.955	0.017*
	(0.004)	(0.005)	(0.006)			(0.005)	(0.009)
Baseline Att >= 0.8	0.957	0.007	0.015**	0.965	0.008	0.948	0.02
	(0.005)	(0.007)	(0.006)	(0.005)	(0.008)	(0.013)	(0.016)
Baseline Att < 0.8	0.958	0.004	0.002	0.959	0.011**	0.957	0.016*
	(0.003)	(0.004)	(0.006)	(0.004)	(0.005)	(0.005)	(0.009)
Panel D: Self Reported Enrollment							
Grades 6-8	0.979	0	0.009*	0.994	-0.004		
	(0.004)	(0.005)	(0.005)	(0.002)	(0.004)		
Grades 9-11	0.989	-0.005	-0.014**			0.988	0.006
	(0.004)	(0.005)	(0.006)			(0.005)	(0.006)
Baseline Att >= 0.8	0.981	0.006	0.008	0.996	-0.008	1	-0.014
	(0.005)	(0.005)	(0.007)	(0.004)	(0.006)	(0.000)	(0.015)
Baseline Att < 0.8	0.983	-0.005	-0.003	0.993	-0.003	0.985	0.010*
	(0.004)	(0.005)	(0.005)	(0.003)	(0.005)	(0.007)	(0.006)
Panel E: Grade 11							
Graduated	0.877	0.029	0.044			0.904	0.041
	(0.025)	(0.043)	(0.031)			(0.024)	(0.034)
Higher Education	0.224	0.049	0.095***			0.197	0.477***
	(0.032)	(0.035)	(0.034)			(0.032)	(0.037)

* significant at 10 percent level, ** at 5 percent level, and *** at 1 percent level
Standard errors are clustered at the school level.

Table 6: Academic Effort, Consumption, and Labor Activities

Outcome Variable	San Christobal Control Average	 Basic - Control	 Savings - Control	Suba, Grades 6-8 Control Average	 Basic - Control	Suba, Grades 9-11 Control Average	 Tertiary - Control
Panel A: Academic Effort, Grades 6-11							
Hours of Homework	2.693	0.015	0.035	2.957	0.025	2.6	0.522***
	(0.033)	(0.044)	(0.048)	(0.040)	(0.040)	(0.072)	(0.116)
Total Grades, Self Reported	0	0.051	0.047	0.024	-0.05	-0.041	-0.051
	(0.024)	(0.034)	(0.036)	(0.034)	(0.058)	(0.040)	(0.052)
Total Grades, Verified	0	0.07	0.045	0.048	0.014	-0.095	-0.061
	(0.034)	(0.058)	(0.047)	(0.048)	(0.055)	(0.056)	(0.085)
Passed in 2005	0.888	0.009	0.020**	0.907	0.011	0.904	0.022
	(0.007)	(0.009)	(0.010)	(0.009)	(0.014)	(0.012)	(0.017)
Panel B: Consumption, Grades 6-11							
Meals Over Last 3 Days	8.018	0.186**	0.244***	8.184	0.056	8.203	0.154*
	(0.051)	(0.079)	(0.073)	(0.059)	(0.086)	(0.074)	(0.080)
Meals with Eggs or Meat	5.07	0.163**	0.184***	5.234	0.037	5.281	0.157*
	(0.042)	(0.076)	(0.053)	(0.050)	(0.066)	(0.062)	(0.093)
Number of Textbooks	0.349	0.038	0.026	0.376	0.031	0.304	0.065
	(0.016)	(0.026)	(0.025)	(0.024)	(0.029)	(0.030)	(0.043)
Panel C: Labor Activities, Grades 6-11							
Total Hours	0.95	-0.344	-0.285	1.006	-0.091		
	(0.166)	(0.253)	(0.230)	(0.157)	(0.174)		
Hours for Pay	0.751	-0.376*	-0.28	0.675	-0.113		
	(0.155)	(0.200)	(0.218)	(0.125)	(0.152)		
Hours without Pay	0.199	0.032	-0.005	0.331	0.023		
	(0.057)	(0.110)	(0.081)	(0.082)	(0.167)		
Panel D: Labor Activities, Grades 9-11							
Total Hours	2.848	0.09	0.386			4.581	-3.022***
	(0.366)	(0.438)	(0.501)			(0.537)	(0.711)
Hours for Pay	2.586	-0.029	-0.003			3.653	-2.246***
	(0.356)	(0.451)	(0.447)			(0.493)	(0.660)
Hours without Pay	0.261	0.118	0.389**			0.928	-0.775***
	(0.089)	(0.101)	(0.153)			(0.204)	(0.215)
Other Children	6.222	0.129	-0.671			5.148	1.677**
Hours for Pay	(0.707)	(1.232)	(0.827)			(0.687)	(0.800)
Other Enrolled Children	0.919	0.415	-0.306			0.869	0.643
Hours for Pay	(0.281)	(0.253)	(0.469)			(0.250)	(0.396)
Other Non-Enrolled Children	5.303	-0.286	-0.366			4.279	1.034
Hours for Pay	(0.644)	(1.213)	(0.766)			(0.646)	(0.896)
Other Children	0.161	0.01	0.118			0.326	-0.117
Hours without Pay	(0.071)	(0.046)	(0.173)			(0.110)	(0.090)

* significant at 10 percent level, ** at 5 percent level, and *** at 1 percent level
Standard errors are clustered at the school level.

Table 7: Heterogeneity of Treatment Effects

Demographic Variable	Control Average	Treatment - Control	Control Average	Treatment - Control	Control Average	Treatment - Control
Panel A: All Students						
	Baseline Attendance		Number of Meals		Textbooks	
All Students	0.779***	0.029***	8.076***	0.181***	0.298***	0.046***
	(0.015)	(0.006)	(0.061)	(0.051)	(0.023)	(0.014)
Baseline Attendance						
Attendance > 0.8	0.823***	0.008	8.022***	0.238***	0.341***	0.017
	(0.014)	(0.009)	(0.080)	(0.082)	(0.034)	(0.032)
Attendance ≤ 0.8	0.760***	0.038***	8.102***	0.145***	0.279***	0.058***
	(0.017)	(0.008)	(0.068)	(0.055)	(0.023)	(0.016)
Gender						
Female	0.792***	0.016**	8.015***	0.170***	0.290***	0.058***
	(0.017)	(0.007)	(0.073)	(0.066)	(0.029)	(0.018)
Male	0.765***	0.043***	8.142***	0.190***	0.307***	0.033
	(0.015)	(0.009)	(0.069)	(0.059)	(0.024)	(0.023)
Income						
Upper Tercile	0.765***	0.049***	8.095***	0.259***	0.319***	0.054**
	(0.019)	(0.011)	(0.084)	(0.085)	(0.031)	(0.021)
Middle Tercile	0.788***	0.026***	8.130***	0.199***	0.285***	0.060**
	(0.015)	(0.009)	(0.067)	(0.070)	(0.032)	(0.029)
Low Tercile	0.784***	0.014	8.002***	0.102	0.290***	0.017
	(0.015)	(0.011)	(0.089)	(0.086)	(0.024)	(0.022)
Household Head's Education						
Upper Tercile	0.770***	0.032***	8.090***	0.228***	0.292***	0.063**
	(0.016)	(0.010)	(0.091)	(0.081)	(0.028)	(0.026)
Middle Tercile	0.783***	0.026***	8.087***	0.173**	0.296***	0.050**
	(0.016)	(0.008)	(0.078)	(0.067)	(0.023)	(0.021)
Low Tercile	0.784***	0.028***	8.045***	0.141	0.309***	0.016
	(0.018)	(0.010)	(0.089)	(0.096)	(0.033)	(0.021)
Panel B: Terciary Treatment						
	Hours Worked		Wage Hours Worked		Wage Hourse, Others	
All Students	3.578***	-3.029***	3.035***	-2.269***	5.770***	1.764**
	(0.334)	(0.709)	(0.311)	(0.666)	(0.546)	(0.836)
Income						
Upper Tercile	3.903***	-4.840***	3.093***	-3.488***	6.323***	4.707**
	(0.712)	(0.980)	(0.599)	(0.753)	(0.912)	(2.001)
Middle Tercile	3.612***	-2.732***	3.165***	-2.392***	4.638***	2.370*
	(0.474)	(0.956)	(0.439)	(0.910)	(0.704)	(1.223)
Low Tercile	3.187***	-0.802	2.829***	-0.675	6.424***	-2.271
	(0.477)	(1.008)	(0.490)	(0.982)	(1.018)	(2.050)
Household Head's Education						
Upper Tercile	2.619***	-1.718*	2.157***	-1.196	4.661***	-0.706
	(0.391)	(0.902)	(0.380)	(0.780)	(0.868)	(1.638)
Middle Tercile	4.248***	-4.096***	3.485***	-2.992**	6.397***	2.348
	(0.570)	(1.384)	(0.464)	(1.192)	(0.908)	(1.902)
Low Tercile	3.855***	-3.464**	3.526***	-3.127**	6.304***	5.091**
	(0.758)	(1.371)	(0.732)	(1.289)	(1.083)	(2.560)

* significant at 10 percent level, ** at 5 percent level, and *** at 1 percent level
Standard errors are clustered at the school level.