

**TESTING ALTERNATIVE APPROACHES TO CONDITIONAL CASH TRANSFER
PROGRAMS IN EDUCATION:
EVIDENCE FROM COLOMBIA¹**

*Preliminary and Incomplete
Please Do Not Cite*

**First Draft: July 2007
Current Draft: Sept 2007**

Felipe Barrera-Osorio (World Bank)
Marianne Bertrand (Chicago GSB)
Leigh L. Linden (Columbia University)
Francisco Perez (Ministry of Education, Colombia)

Abstract: Using an RCT design, 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). More generally, the results suggest that credit constraints may not be involved in decision to attend secondary school, but may be involved in the decision to attend a tertiary institution. Thus, it may be more effective to focus on completing academic goals (such as graduating or tertiary enrollment) rather than attendance in secondary school. Finally, our strategy allows us to assess the effects of the treatment on students' families and peers. We find that the subsidies cause a reallocation of responsibilities within the household with treated children attending school more and working less while untreated household members do the opposite. We also find that indirect peer influences are ale relatively strong with indirect effects that are on average a quarter of the direct effects.

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. Fedesarrollo, the think thank for which Barrera-Osorio and Perez were working at the execution of the project, provided financial support as well and helped the SED in the design and implementation of the program. 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 the logistical assistance and for the data collection. 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.

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). In the long-term, there allowing children to attend private schools improves completion rates (Angrist et al, 2006; Angrist et al., 2007). However in the short-term, improving quality does not seem to be a major inducement. Interventions proven to improve the quality of education generate no changes in participation levels (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 (Schultz, 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 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

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).

savings treatment, but also pays a large subsidy that incentivizes 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 are also able to test the relative importance of two methodological improvements over previous education-based conditional cash transfer studies. First, we 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. Second, we directly map students' friendship networks in order to compare the performance treatment estimates using these explicit maps to those that use the indirect measure of peer effects measured through the treatment density of various cohorts to which a child belongs.

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

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

institutions. The tertiary treatment generates 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. We observe a small decrease in the number of hours worked overall, and a much larger reduction for the Tertiary Treatment. Not surprisingly, the effects are the strongest for children who probably would not have met the attendance targets without the incentive, and we see no effect on attendance for the poorest families.

Externalities also play an important role. Within families, a larger fraction of treated school-aged children causes a reallocation of academic opportunities and labor market responsibilities within the family. Untreated individuals who are not already focused on work change their focus from academics and household activities to work. For school-aged children this means lower rates of participation in school with 3 percent less students attending school for every 10 percent of the school-aged family members treated. Similarly, within peer networks, more treated friends encourage higher attendance rates equivalent, on average, to a quarter of the direct effect of the treatment.

Finally, the methodological improvements that we implement are also significant. Self-reported attendance rates are significantly higher than those measured through direct verification. And we estimate much lower attendance effects for students who completed our baseline and follow-up surveys, suggesting that the repeated visits to schools may be necessary to capture the effects of the program on those most likely to need and respond to the program. Finally, we are also able to demonstrate that explicit mapping of friendship networks provides a more accurate estimate of peer effects than the indirect effects often estimated by measure the fraction of various cohorts receiving a treatment.

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 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 600,000⁵ 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

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.

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.

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.

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 Cristobal and Suba. Eligible registrants in San Cristobal 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 Cristobal 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.⁶

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

⁶ The over subscription and recruitment process are based on the techniques used in the assignment of school vouchers in the PACES program implemented nationally in Colombia. This process is described in Angrist et al. (2002).

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.⁷ 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 Cristobal 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.⁸ This left 17,309 students eligible for the two experiments: 10, 947 in San Cristobal 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

⁷ 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.

⁸ Unfortunately, this program was never actually implemented, and this portion of the study was discontinued due the lack of an intervention.

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 in San Cristobal was about 63 percent while in Suba the 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.⁹

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.

⁹ 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.

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.

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:

$$x_{ij} = \beta_0 + \beta_1 Treat_i + \varepsilon_{ij} \quad (1)$$

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 Cristobal:

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

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.

In addition to the direct estimates of the programs, we also estimate the external effects of the treatment on students' family members and peers. For these specifications we are interested in the relationship between the individuals' behavior and either the fraction of peers treated or the fraction of school-aged family members treated. To do this, we have to account for the fact that the fraction of registered peers or family members is possibly endogenous. As a result, we use an instrumental regression model in which the fraction of treated peers or school-aged family members is instrumented with the fraction of registered peers or family members who receive the treatment. For the friendship networks, the specification takes the following form:

$$y_{ij} = \beta_o + \beta_1 Frac_Treat + \beta_2 Frac_Treat^2 + \beta_3 Treat + \delta X_{ijk} + \phi_j + \varepsilon_{ij} \quad (3)$$

All of the variables are defined as before and β_1 and β_2 are the estimated effects of the fraction of friends treated by the program. For non-treated members of the family, we use a similar specification except that we omit the school fixed-effects and cluster the standard errors at the family level.

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 use 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:

$$y_{ijk} = \beta_o + \delta X_{ijk} + \varepsilon_{ij} \quad (4)$$

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 is 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 Cristobal, 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 assess 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 Cristobal (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 none¹⁰ 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

¹⁰ 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.

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 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 Cristobal 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 Cristobal, 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 Cristobal 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.¹¹ 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.

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

¹¹ We obviously have concerns about the accuracy of self-reported enrollment rates, given the results for the attendance rates.

(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, breaking out grade 11 when we focus on school expenses. Panel A displays the results for our measures of academic effort. Panel B contains the results for consumption variable. And finally, Panel C contains information on labor market activities. The columns are defined as in Table 5 and all estimates are again done using Equation 2.

In general, the treatments have little effect on our measures of academic effort in Panel A. Only two effects are statistically significant. First, 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. Second, the Savings Treatment causes a 2 percent increase in the number of children passing in the last year. The other treatments have a similar effect (1-2 percent), but none of them are statistically significant.

Finally, looking to students grades, it would be surprising to see changes in grades as a result of only increased exposure to school (given the magnitude of the changes in attendance), 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.¹²

Panel B show the results for consumption and is divided into food and school expenses in the first and second two rows respectively. The effects on food consumption are small but

¹² It is interesting to note the relative similarity of the effects estimated through self-reported and verified grades. Generally, students tend to slightly overestimate their grades, but the treatment and control students seem to overestimate to the same degree, resulting in similar estimates of the programs effects.

largely uniform across the treatments. 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.

The second two rows of Panel B contain the results for school expenses. This variable is the sum of questions that require the family to provide details (in thousands of Colombian pesos) of their educational expenses for the entire year, including administrative fees, school supplies, transportation, and ancillary expenses. Consistent with the relative effects of the Basic, Savings, and Tertiary Treatments, the treatments seem to have little effect on spending in grades 6-10, suggesting that financial constraints are not significant for individuals in these grades. However, the Tertiary Treatment seems to have a large effect on educational expenditures for individuals who are of age to attend a tertiary institution. Not all of the 600,000 pesos are spent on academic activities. Families would have spent 140,000 pesos anyway on average, but the treatment causes families to spend an extra 240,538 pesos. In total, it seems that 63.4 percent of the subsidy is spent on academic expenses.

Finally, Panel C contains the results for labor market activities. The first three rows show the average for an indicator variable listing the primary activity of the child (studying, work, or taking care of the household). Finally, the last two rows show the number of hours worked in the last week that the child worked and the amount of money earned during that week (in 1,000's of pesos). Again, the basic and savings treatments show little effect, while the Tertiary Treatment shows significant effects. It increases the the percentage of kids whose primary activity is attending school by 12 percentage points, and reduces by 4 percent points the fraction of kids whose primary activity is working or taking care of the home. Consistent with these changes, children receiving the Tertiary Treatment work an average of 2 hours less a week and earn 2,000 pesos less a week.¹³

3. Effect on other Household Members

If families redistribute academic opportunities and labor market responsibilities to reflect the relative costs of a child's education, then the receipt of the treatment by one child may affect the

¹³ These results are consistent both for children who were in grade 11 in the last year and who were in grades 9 and 10.

activities of other children. To investigate this, we estimate the effect of receiving a subsidy on non-treated school-aged children and adults in the family. For each variable, we use the instrumental variables model described Equation 3 to estimate the relationship between labor market activities and the percentage of school-aged children treated in the family, using the fraction of registered children assigned to each treatment as an instrument.

The results are presented in Table 7. They are grouped by age and gender, presenting the results for males and females of working (between 24 and 60) and school-age (7 and 24).¹⁴ For school-aged children, we report primary activities, school participation, and the number of hours worked for money in the last week the child worked. For working-age family members, we report only primary activity and the number of hours worked for wages in the last week worked. The columns are arranged as in Tables 2-6, with the control average estimated from families who do not receive a subsidy.

To interpret these results, it is important to understand the structure of the families in the study. The families are surprisingly homogenous with between 1 and 3 children. On average, each family has 2.5 children between the ages of 7 and 23, but 83 percent of families have 3 children or less. Since most families receive one subsidy (only 9 percent of families receive 2 or more), a families receipt of a subsidy changes the percentage of children treated by 40 percent on average, and this percentage varies between one third and one hundred percent if the family is treated. However, because most families only receive one subsidy, the variation in treatment density is generated from variation in family size – higher treatment densities reflect the effect experienced by smaller families.

Turning to Panel A, we see that the effects, on school-aged boys, of a family receiving a subsidy are large and consistent across the various treatments. Children not receiving the subsidy study less, work more, and have completed few grades after being exposed to the treatment for only a year. Without the treatment, 75-80 percent of children study as their primary activity and 13 to 17 percent work (very few do household work). For every 10 percent of the school-aged children who receive the treatment, about 3 percent less children focus on studying and switch to work. Given the average change in treatment percentage of 40 percent, the treatments caused an average of 12 percent of children to change their focus from studying to

¹⁴ We chose the age range for school-aged children by using the range of ages for children who report being enrolled in school in the SISBEN survey. The results are consistent for other choices of the age range.

work. A similar number of children report not being enrolled in school (from about 80 percent), and children report having completed half to eight tenths of a grade. Consistently, boys then work a bit more than an hour a week for every ten percent of the children treated for an average effect of about 4.5 hours more.

Turning to school-aged girls in Panel B, the results are similar with two exceptions. First, while a similar percentage of girls no longer report studying as their primary activity, their focus changes not just to work but also to household assistance. A ten percent increase the fraction of the children who are treated increases the fraction of school-aged girls working by 2 percent and taking care of the home by 1 percent. Correspondingly, girls report an increase in the number of hours worked that is 60 percent less than the effect for school-aged boys.

The results for working-aged family members presented in Panels C and D show a similar difference in response by gender. Panel C contains the results for men. On average, 90-94 percent of men report working as their primary activity while about one percent of them report assisting with the household as their primary activity. As a result, men in families with no treatments work about 45-47 hours a week, and generally, men show no response to an increase in the fraction of children receiving a subsidy.

Women in Panel D, on the other hand, show an almost even initial distribution of primary activity between work (50-55 percent) and household activities (40 percent). The point estimates suggest a slight shift from household activities to work as a result of the subsidies, but the results are not significant. However, for the basic and savings treatment, the amount that working-aged women work as a result of the program increases by a statistically significant 0.35 to 0.41 hours per tenth of the children treated.

These results suggest that the family reallocates both academic opportunities and labor market responsibilities as a result of the subsidy. Even though, in most cases, the immediate activities of the treated children fall by less than a half an hour a week (except for the Tertiary Treatment), the rest of the family reallocates effort away from school and towards alternative activities. Only working-aged males, who are already completely focused on labor market activities, do not respond to the family receiving one of the subsidies. This discontinuous response seems to be an over-reaction to the change in the short-term costs of sending an individual child to school, but may reflect an (unobserved) change in the longer-term anticipated career-path of the children.

Finally, while it is tempting to interpret these results as negatively impacting the welfare of non-treated members of the family, the available data is insufficient to support this conclusion. By providing subsidies to the families, the treatment both changes the relative costs of sending children to school and provides the family with resources that it would not otherwise have. It may be that the family allocates resources efficiently and uses these funds to increase the long-term welfare of every member of the household. This might occur, for example, through future transfers of resources from more educated children to less educated members of the family.

4. Heterogeneity in the Treatment Effects

Another important dimension of these incentive programs is their relative impact on different types of students and families which we investigate in Table 8. In Panel A, we study possible heterogeneity in the effect on treated students 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 hours worked. However, in order to maintain sufficient sample size within subsets of the data, we pool the data and estimate the average effects across all three treatments.

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 as a reference. It is important to note, however, that when we pool the results from all of the treatments, many of the effects that were small and insignificant are now statistically significant because of the increased precision of the estimates. So, for example, the treatments seem to have an average effect of reducing the number of hours students worked by a little less than half an hour, a difference that is now significant.

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 have already seen, only students whose attendance without the program would be under 80 percent increased participation as a result of the program. 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.

The cash transfers do seem more generally to 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 demonstrates that outside of the incentive effects, the small transfers from the program can have a measurable effect on general wellbeing. Surprisingly, as shown in columns 5 and 6, the estimated difference in hours worked also does not depend on the projected attendance rates of the children.

The next two rows divide the sample by gender. There seems to be a difference in the responses of boys' and girls' attendance patterns (4.3 percentage points to 1.6 percentage points).¹⁵ However, there seems to be no difference in the number of meals or number of hours worked by boys or girls.

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). The same relative effects are observed for the number of meals and the number of hours worked by students. These results suggest that the transfer may just be too small to make a difference in the lives of the poorest families. While all of these students are poor, the poorest of our sample are extremely poor (reporting a total income of less than 3 USD a day). Apparently, those who are less poor have the capacity to take advantage of this program while those with fewer resources do not.

Finally, Panel B estimates the relative effects for other school-aged members of the students' households (boys and girls). The layout of the table is different from Panel A in that it focuses on household rather than individual characteristics. Heterogeneity in effects are

¹⁵ 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.

provided for the probability that a child's primary activity is studying, self-reported enrollment, and the number of hours worked during the last work week.

Two patterns are evident in these results. First, as observed in Panel A, the effects are strongest for the relatively wealthiest families. All children show a response to having a child in the family treated, but the poorest families experience the smallest effects. Second, the effect is also strongest for families whose heads are the least educated (controlling for income). This may reflect a lower value for general education in these families.¹⁶

5. Peer Effects

Finally, because our randomization strategy induces intra-school variation in the treatments, we can directly measure the effects of students on each other through our explicit mapping of friendship networks collected in the baseline survey. Such externalities have been previously estimated for conditional cash transfers by exploiting the fact that the original PROGRESSA program only provided subsidies for the poorest families in a village. Both Bobonis and Finan (2006) and Lalive and Cattaneo (2006) compare the school attendance patterns of children whose families were too rich to receive the subsidy between treated and untreated villages. This comparison allows them to experimentally identify the effects of the externality, but while they estimate the effects as a function of the fraction of treated and untreated children, the fraction of treated children is not experimentally assigned. So, the design does not allow for the estimation of the causal effects of changes in the treatment density of a network.

Our design provides both an explicit mapping of children's friendship networks and random assignment of the fraction of a student's friends who are treated. This allows us to compare directly students who have more treated friends to those who have less using Equation 3. As in the family regressions, the possible endogeneity arising from the fact that registration is likely to have followed friendship ties is controlled by instrumenting for the fraction of registered friends actually receiving the treatment. These results are presented in Panel A of Table 9. In Panel B, we replicate the specifications used in Finan and Babonis (2006) and Lalive and Cattaneo (2006) by estimating changes in student behavior as a function of the fraction of

¹⁶ In unpublished results, we divided the students in Panel A in the same way, but the results were not consistent with those in Panel B with the exception of the number of hours worked.

students treated within school-grade-gender groups. Panel B thus illustrates the difference between the direct and indirect mapping of friendship networks.

Starting in Panel A, we show the results of four regressions. Unlike the previous tables, each column represents a single regression of the specified outcome variable on the individual dependent variables. The first column provides the reduced form regression of verified attendance on the fraction of registered friends who were treated. The second column estimates the first stage regression of the fraction of a student's friends who are treated on the fraction of registered friends who are treated. (As expected there is a significant amount of correlation between these two variables with almost 70 percent of the variation in treatment density explained.) The third column provides the OLS regression, and finally, the IV estimate is provided in the fourth.

The reduce form, OLS and IV estimates all show a quadratic relationship between students' attendance and the fraction of students' friends who are treated. There seems to be, however, a negative bias to the registration of peers because the IV estimate increases the linear and quadratic effects of the OLS estimates from 0.045 to 0.173 and from -0.069 to -0.389 respectively. This suggests that the treatment of friends has a positive effect on a students attendance rates, but that this effect declines as more friends are treated. On average, the implied indirect effect is 0.19 percentage points or about a quarter of the direct effect of 0.76 percentage points.

Panel B shows the results of our replication effort. The panel is organized like Panel A except for the fact that the treatment effect is estimated as a function of the fraction students treated within school-grade-gender cohorts. While the first stage regression again seems strong, the point estimates for the regressions in columns 1 ,3, and 4 are very small and insignificant. The density of the treatment within grades does not seem to strongly affect the attendance rates of students. This is a striking contrast to the results utilizing the explicit friendship mappings. The differences suggest that such indirect measures of externalities tend to understate the true effects.

V. Conclusion

This project demonstrates that experiments involving conditional cash transfer programs can be used to understand how variation in the structure of these programs may affect educational and related outcomes for targeted children and their families. In general, 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. 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. And researchers need to focus attention the response of families, rather than just individual children, to education programs.

Specifically, 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.

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.

Finally, externalities are extremely important. Within families, we observe a general reallocation of responsibilities that suggest families allocate academic opportunities and labor market responsibilities as a function of the relative costs of children's education. We observe that treated children increase their academic involvement and reduce the time spent at work, while that untreated family members who are not already completely focused on work, reduce their participation in non-work related activities (education and household chores) and increase their labor market participation. On average, untreated siblings, work 3-4 more hours a week and are 12 percent less likely to be enrolled in school. Similarly, peers can also affect students'

behavior. We estimate that peer-related externalities can influence behavior by almost a quarter of the direct program effect.

References

- Angrist, J. D., E. Bettinger, E. Bloom, E. M. King and M. Kremer. (2002) "Vouchers for Private Schooling in Colombia: Evidence from a Randomized Natural Experiment," *American Economic Review* 92: 1535-59
- Angrist, Joshua D., Eric Bettinger, and Michael Kremer. (2006) "Long-Term Educational Consequences of Secondary School Vouchers: Evidence from Administrative Records in Colombia." *American Economic Review* 96:847-862.
- 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.
- Banerjee, A., S. Cole, E. Duflo and L. Linden (2007) "Remedying Education: Evidence from Two Randomized Experiments in India," Forthcoming *Quarterly Journal of Economics*.
- 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 an Unusual Policy Experiment," *American Economic Review*. 91(4): 795-813.
- Fedesarrollo (2005) "Proyecto 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.
- He, F., L. Linden and M. MacLeod (2007) "Teaching What Teachers Don't Know: An Assessment of the Pratham English Language Program" Columbia University Department of Economics *Mimeo*.
- 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*.
- Pritchett, L. (2004) "Towards A New Consensus for Addressing the Global Challenge of the Lack of Education" Copenhagen Consensus Challenge Paper in Education.
- 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) *Development Goals Report 2006*, United Nations, New York.
- Vélez, C.E., E. Castaño and R. Deutch (1999) "An Economic Interpretation of Colombia's SISBEN: A Composite Welfare Index Derived from the Optimal Scaling Algorithm" Mimeo, Poverty and Inequality Advisory Unit, Inter American Development Bank, Washington D.C.
- Vermeersch, Christel and Michael Kremer, "School Meals, Educational Achievement, and School Competition: Evidence from a Randomized Evaluation" World Bank Policy Research 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 2006*, The World Bank, Washington.

Figure 1: Distribution of Family Income




Figure 2: Distribution of Attendance

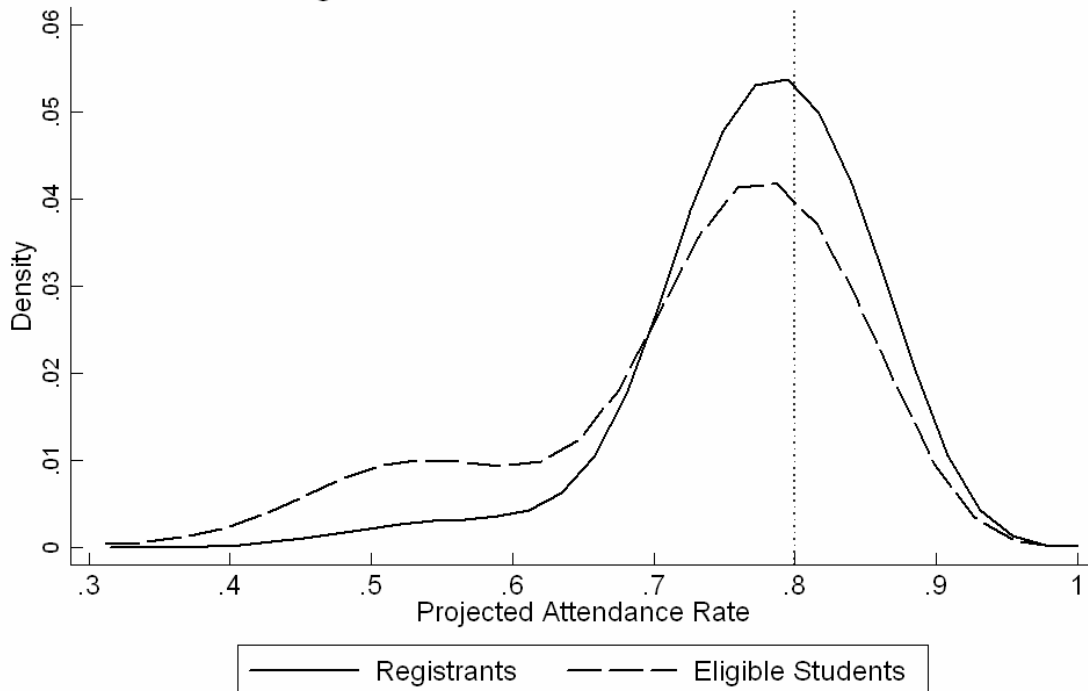


Figure 3: Distribution by Family Income at Baseline




Figure 4: Distribution of Attendance at Baseline




Figure 5: Distribution of Attendance at Follow-Up





Figure 6: Actual Attendance by Predicted Attendance



Note: Results from local polynomial regressions (bandwidth=0.075)