Using Conditional Transfers in Education to Investigate Intra Family Decisions: Evidence from a Randomized Experiment[1]

First Draft: July 2007
Current Draft: Dec 2007

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

Abstract: We evaluate multiple variants of a commonly used intervention to boost education in developing countries – the conditional cash transfer (CCT) – with a student level randomization that allows us to generate intra-family and peer-network variation. We test three treatments: a basic CCT treatment based on school attendance, a savings treatment that postpones 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. On average, the combined incentives increase attendance, pass rates, enrollment, graduation rates, and matriculation to tertiary institutions. Changing the timing of the payments does not change attendance rates relative to the basic treatment but does significantly increase enrollment rates in at both the secondary and tertiary levels. Incentives for graduation and matriculation are particularly effective, increasing attendance and enrollment and secondary and tertiary levels at a higher rate than the basic treatment. We also find that indirect peer influences are relatively strong with the average magnitude similar to that of the direct effect. The size of the effect increases with the number of friends treated, but at a declining rate.

[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 and graduation rates. The challenge, however, is getting the kids in school. 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 with older children, 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 average quality of education is substantially lower (Pritchett, 2004). 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, interventions that directly change the cost of attending school do seem to work. Families, for example, respond to direct inducements to attend such as school meals or direct cash incentives (Vermeersch and Kremer, 2005; Schultz 2004; Glewwe and Olinto, 2004; Schady and Araujo, 2006; Schultz, 2004; Attanasio et al, 2006 among others). Families also respond to direct reductions in the cost of education either through subsidies to attend private schools (Angrist et al 2002, 2006), reduced user-fees (Barrera, Linden, Urquiola, 2007) or through scholarships (Kremer, Miguel, and Thornton, 2007).[2]

Given the apparent importance of family finances, it may be possible to improve the performance of direct incentives and better understand their effects by taking into account recent research into the family decision processes surrounding the allocation of resources within the household. As an example, we focus on two general themes of this literature – the difficulties families face when saving (for example, Ashraf, Karlan, and Yin, 2006) and the unequal allocation of resources within the household (for example, Blundell, Chiappori, and Meghir, 2005). Our strategy is to build on this existing research by testing multiple variants of a single, well-established direct incentive program: the conditional cash transfer. To date, most programs offer an incentive similar to that used in the original Mexican PROGRESSA program that pays students on a monthly or bi-monthly basis for meeting a specified attendance (usually 80 percent per month) or enrollment target. We build on this basic model by evaluating multiple variants of conditional cash transfer treatments in a large municipality: Bogota, Colombia. Our research strategy allows us to relax families' savings constraints for school fees and to incentivize directly graduation and matriculation to tertiary institutions. It also generates exogenous variation in treatment within family and friendship networks, thereby allowing us to directly identify the influence of siblings and peers through the family and peer networks.

Specifically, we first use the basic treatment implemented in a manner very similar to the original PROGRESSA program. The effect of this basic treatment is compared to that of a second intervention ("the savings treatment"). Using the same conditions, the savings treatment 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 second intervention 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. We also test, in a second experiment, a third treatment, called "tertiary treatment." This third treatment provides children with the

---

[2] Health is also a major factor in determining school attendance (Miguel and Kremer, 2006; Bobonis, Miguel, and Sharma, 2006).

[4] One exception is the system that the national Colombian government used to allocate school vouchers in the Program de Ampliacion de Cobertura de la Educacion Secunderia (PACES) program (see Angrist et al, 2002 and Angrist et al 2006). While our use of this allocation strategy was a practical solution for conducting a student-level randomization, the intra-family and intra-friendship network variation enabled by this strategy have direct policy implications for this an other allocation mechanisms that partially families or groups of friends.

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 most previous studies.[4] 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.

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 avoid the self-reporting bias from the survey data used in most 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, allowing us to assess directly the influences of peers on students' attendance rates.

Taken together, all of the incentives generate significant changes in the behavior of students directly treated by the program. Students are more likely to attend school (2.8 percentage points), more likely to remain enrolled (2.5 percentage points), more likely to matriculate to the next grade (1.6 percentage points), more likely to graduate (3.9 percentage points), and more likely to matriculate to a tertiary institution (21 percentage points). For daily attendance, the effect is much stronger for students who would not have met the attendance target without the program.

Within the treatments, however, the form of the incentive matters significantly. Simply changing the timing of the transfer with the savings incentive increases enrollment in both secondary and tertiary institutions over the basic treatment (by 3.6 and 3.3 percentage points respectively) while not reducing the daily attendance rates of students despite the lower monthly transfers. Compared to the basic treatment, the tertiary treatment encourages higher levels daily attendance (3.5 percentage points more for students least likely to attend) and higher levels of enrollment at the secondary (3.3 percentage points) and tertiary levels (46 percentage points).

We also observe important spillover effects of the program within. Treating friends encourages higher attendance at a similar magnitude as the direct effect. This effect appears to be independent of students' direct receipt of the incentive or of a students' counter-factual attendance rates. However, the gains in attendance decline sharply in the fraction of friends treated – one treated friend has significant benefits, but an additional friend has almost no additional effect.

Finally, our results have several methodological implications. Self-reported attendance rates are significantly higher than those measured through direct verification. Second, we estimate much lower attendance effects for students who completed our baseline surveys, suggesting that those most likely to need and respond to incentive programs may also be the most difficult to observe directly. 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 simply measuring 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. Section 4 presents verification for external validity, balance of the baseline and measures of attrition. We present the results of the analytical models in Section 5. Finally, we conclude in Section 6.

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, with only 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 as well as of 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 close to 90 percent of these funds are required by law to go toward health and education. With these funds, municipalities must develop, maintain, and run the facilities in their jurisdictions. Municipalities that have greater capacity to collect and administer taxes supplement central resources with local resources, usually from property taxes.

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 children are legally 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 old, 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 in the city of Bogota. 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. 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. Based on the responses to our surveys, the total annual value of the transfer (300,000 pesos) is three times more than what students report earning on average and is slightly more than the average 250,000 pesos that families report spending each year on educational expenses. 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.[5] 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$10) and the remaining third was held in account. The accumulated funds were then made available to students families during the period in which students enroll and prepare for the next school year. If students reached the attendance target every month, this treatment would make 100,000 pesos (US$ 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, 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 (Tertiary 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 in the Savings 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

[5] 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. However, the inflation rate in 2005 was only 4.85%, and the net effect of this difference is to reduce the value of the savings treatment which, we will show, is more effective than the basic treatment despite the slightly lower value.

earn the right to receive a transfer of 600,000 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. It is important to note, however, that unlike the savings treatment, the tertiary treatment does more than just change the timing of the payments to families because the total value of the tertiary treatment for students in each grade is greater than the equivalent value of the basic treatment.

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 were 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 research design 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.[7]

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. [8] 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.[9] 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

[7] The over subscription and recruitment process are similar to 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).

[8] 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.

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

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, it, of course differs between them. The probability of treatment is 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.

C. Data

The richness of the available data is one of the major strengths of our study. The data come from six 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 subsets of 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 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.[10]

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.

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.

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.

[10] 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.

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 the 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.

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 families.

Finally, we obtained administrative records from the Secretary of Education that includes the enrollment records of every child in a public school and many private schools in the two localities. This data allows us to assess the effect of the treatments using every student that registered for the randomization, including those not claiming to attend one of the 68 schools selected for attendance data collection and surveying. Combined with the other outcome variables, this provides us with three concentric groups of students: all of the students who registered for the randomization (for which we only have administrative enrollment data), all of the students registered at the 68 schools selected for surveying (for which we have both administrative enrollment data and verified attendance data), and those students at the 68 schools who completed baseline and follow-up surveys (for whom we have administrative enrollment data, verified attendance data, and the information collected in the surveys).

D. 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. Third, we use an instrumental variables model to estimate externalities generated by the treatments within families and students friendship networks. 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, in order to validate the randomization, 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 Cristobal:

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

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.

In addition to the direct estimates of the programs, we also estimate the external effects of the treatment of students' peers. For these specifications we are interested in the relationship between the individuals' behavior and either the fraction of peers 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 variables regression model in which the fraction of treated peers is instrumented with the fraction of registered peers who receive the treatment. The specification takes the following form:

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

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.

Finally, we use one last specification to estimate what the attendance and enrollment rates of students who received the treatment would have been without the treatment.[11] We estimate these counter-factuals by modeling the behavior or students in our control groups using only the available baseline demographic characteristics. For treated students, we use their baseline characteristics and the coefficients from our regressions on the control group to project what these students' would have done had they not been treated.

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

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

[11] 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.

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, and it contains no new information.

IV. External validity, baseline balance and attrition

We proceed as follows in this section. 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 problems of randomized evaluations is that, because they often focus on specific group of individuals, 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 registration for 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 20 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 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.

Finally, because students are eligible for the Tertiary Treatment only in Suba, we need to make sure that the students in Suba are 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 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 statistically significant (such as the fact that the Basic Treatment has 3 percent more girls than the savings treatment) 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. An example is shown in Figure 3. The figure contains a plot of the distribution of household income in the treatment and control groups – as shown in the plot, the distributions are very similar.

C. Attrition from Baseline

Comparability at baseline is critical, but even if the two groups are comparable at baseline, 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 attritors 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 all 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 to 7 years difference), the age of children in San Cristobal (2-3 years), and the years of education of students in Suba grades 9-11 (1.25 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.

V. Results

A. 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. It is important to note, however, that within secondary school, the treatments create separate incentives for attendance and enrollment. Within the academic year, the attendance targets encourage students to attend school regularly. Between academic years, however, the anticipated value of the transfers in the following year, encourage families to re-enroll children in school.

Because the program started during the 2005 academic year and required that all students already be enrolled, we first consider students attendance rates as measure by our team of attendance monitors during the last few months of the 2005 academic year. The overall average effects of the treatments combined was to increase verified attendance at school by 2.8 percentage points which is statistically significant at the one percent level. The pooled effects of the treatment are graphically depicted in Figure 4 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[12] 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, which is presented in Figure 5. 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). The effect is fairly uniformly distributed across families. Families who would have met the 80 percent target without the intervention seem to respond slightly less strongly than those who would not have met the target, but interestingly the program seems to make a difference even for students who would not seem to have a strong incentive to increase their attendance rates. Figure 6 performs the same exercise using the administrative enrollment measure, but here the implications of the incentives are significant. Unlike Figure 5, students who would have been likely to re-enroll show no effect from the program while those with a probability of re-enrolling of less than 80 percent show an effect that is fairly evenly distributed across students.

Dividing up these effects to test separately for the individual effects of each treatment, we turn to Table 5. Panel A contains the results for our first outcome variable – verified attendance for all students in the 68 schools selected for surveying during the end of 2005. Panel B contains

[12] 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.

the same verified attendance measure but just for those students who were found in the follow-up survey. Panel C contains the administrative enrollment rates for each subset of our sample, and D contain self reported attendance rates for 2006 based on the follow-up surveys. And finally, Panel E contains the variables pertaining to students in Grade 11 who were in their terminal year of secondary school in 2005. 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. Finally, the last column contains the overall average effects by pooling the effects of the individual treatment. 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 the average results for all students. The second row contains students in grades 6-8 and the third row contains grades 9-11. Overall the basic and savings treatment increase attendance in San Cristobal by 3.3 and 2.8 percentage points respectively. The results are evenly distributed across the grades. For grades 6-8 in San Cristobal, the Savings Treatment increases attendance by a similar amount as the Basic Treatment (3.5 and 2.6 percentage points respectively), despite the lower monthly transfer. Interestingly, the Basic Transfer has no effect on attendance in Suba, though if we pool the sample the estimated effect is a statistically significant 1.9 percentage points. And although the point estimate is then 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 similar, but the results for the Suba experiment (the Tertiary Treatment) are an increase in attendance by 5.0 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 shown in Figure 5, the treatment was slightly more effective for students whose projected baseline attendance was below the attendance target – a difference of about 1 percentage point, although this difference is not statistically significant at conventional significance levels. Only the basic treatment in Suba does not follow this trend, revealing a higher attendance effect for students likely to attend without receiving the treatment.

In Panel B, we focus on just those individuals who were found at school in the baseline survey. From the first three rows, it is clear that primary drivers of the effects reported in Panel A are the students who were not found in the baseline survey. Despite the fact that we returned multiple times to administer the baseline survey to students who may have been absent form the initial visits, individuals found in the baseline survey show a much smaller response to the treatment. 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 were found in our baseline survey are different than the average students that we selected from the registration process. Looking at row 5, even conditioning on our baseline attendance measure being less than 80 percent, the students whose baseline characteristics fit this classification attended an average of 83 percent of the time. So, conditional on observable characteristics, those students in our baseline and follow-up surveys attend school much more often. Second, in this subsample, the main difference is again between high and low attending students, but the difference is starker than in Panel A. In San Cristobal, those projected to attend less than 80 percent of the time show an increase in attendance of 1.8 percent due to the Basic Treatment and of 2.2 percent due to the Savings Treatment. In Suba, there was no response to the Basic Treatment, but the Tertiary Treatment increased attendance by a statistically insignificant 2.8 percentage points. In this sample, the differences between students with high and low baseline projected attendance are statistically significant for the basic and savings treatment (p-value of 0.004) and the regression pooling all treatment (p-value of 0.087).

Turning to Panel C, we estimate the effects of each treatment on enrollment in secondary school in 2006 using the administrative data provided by the Secretary of Education. The first row contains the results using all available students. Both the savings and tertiary treatments significantly increase enrollment by 3.6 and 3.3 percentage points respectively while the basic treatment has no effect. We then divide the sample first by grade and then by whether or not students were likely to attend school without the treatments. The basic treatment has no effect on any of these subsamples. The savings treatment, however, generates sizable effects in all samples except for those students likely to re-enroll even without the subsidy, and the largest effect is for students who are least likely to reenroll without the treatment. Unlike the attendance effects, these differences between the basic and savings differences are statistically significant in the regressions using all students (p-value of 0.024), students in grades 6-8 (p-value of 0.042), and students unlikely to attend without the subsidy (p-value of 0.005). The difference for students in grades 9 and 10 is almost significant (p-value of 0.12) at conventional levels.

In Suba, the results are different. Here the basic treatment does have an effect on students in grades 6-8 of 2.7 percentage points, a difference which again is driven by the re-enrollment of students least likely to enroll without the subsidy. The basic treatment has an effect of 4.9 percentage points on these students. For older students, the tertiary treatment again has a very large effect. Overall it causes a 3.3 percentage point increase in enrollment rates and for students unlikely to attend without the subsidy, it increase enrollment by 8.0 percentage points. These results suggest that incentivizing graduation and tertiary enrollment can be a powerful incentive for students' academic participation.

Next, we can estimate the effects on enrollment using the same sample of students we used in Panel B, the students who gave a baseline and followup survey. As with the attendance results, the effects of the program are much smaller for these students, and in this case, none of the treatments seems to have had an effect on these students' enrollment rates. The Suba treatments show differences of 2.7 percentage points, but neither of these is statistically significant. Finally, we can also compare these estimate effects to our self-reported attendance rates provided in row 4. While the self-reported estimates also show no effect, it is clear form the average rates of the control group that the self-reported attendance rates are significantly biased upwards – so much so that if an effect existed, it would be difficult to detect.

Similarly, we also collected attendance data for student in 2006. This information is displayed in Panel D. Since students are coded as having a zero attendance rate if they report not attending school in 2006, these estimates should at least reflect the differences in enrollment show in Panel C, but they do not. Like the self-reported enrollment rates, the average responses from all groups seem to be over-estimated, and these overestimated participation rates do not reflect the differences observed in the administrative and directly monitored data.

Panel E contains the results for students who were in the 11th grade in 2005 and should have graduated at the end of the year. While this information is self-reported like the attendance and enrollment data in Panels C and D, these estimates seem to be much more realistic. The first row contains an estimate of whether students report having graduated – most but not all students (88 and 90 percent) do in fact report graduating. The point estimates on the treatment effects for each treatment are positive and very similar in magnitude, but insignificant, possibly due to the significantly smaller sample than was available for the previous estimates. However, polling across treatments, the results do have an average effect of 3.9 percentage points that is significant at the 10 percent level.

Looking at whether or not students report both that they have graduated and continue to attend an educational institution, 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 49.7 percentage points is dramatically higher than the effect of a statistically insignificant 4.0 percentage points for the Basic Treatment. The Savings Treatment also increased the enrollment rate by more than the Basic Treatment (8.8 percentage points), but the difference between the two treatment effects is again not statistically significant.

Despite the numerous studies of programs designed to keep students enrolled in school, comparisons to other research is complicated by differences in methodology and the target populations of the programs. For example, most conditional cash transfer programs provide non-education forms of assistance along with educational incentives (usually for nutrition or health care) and are targeted at more rural populations of students in primary and lower secondary school (usually grades 1-8). In addition, the majority of studies employ non-experimental evaluation strategies and focus exclusively on primary and secondary enrollment. We know of no other study that measures the impact of incentive on matriculation to tertiary institutions or that is able to directly verify students' attendance patterns.

Despite these differences, our estimates of the impact of the savings and tertiary treatment on secondary enrollment are a bit larger but of the same magnitude as those measured in other contexts. For example, within the context of the experimental evaluations, our estimates are comparable to those in Honduras (rural students aged 6-13, 1-2 percentage point increase in enrollment, Glewwe and Olinto, 2004), Ecuador (rural students aged 6-17, 3.5 percentage point increase, Schady and Araujo, 2006), and Mexico (rural students grades 1-9, boys 2.5 percentage points, and girls 3.5 percentage points, Schultz 2004). Attanasio et al (2006) assess, with propensity score matching, a similar program to ours (but that includes a nutrition component) run by the national Colombian government and find an in crease in self-reported enrollment of 4.8 percentage points for urban children aged 14-17. However, these areas are much less developed than Bogota, and students have a baseline enrollment rate of 65 percent.

The closest study to ours both in methodology and target population is the evaluation of the PACES program that provided vouchers to students entering the sixth grade. Over three years, Angrist et al (2002) find an increase in the percentage of students who completed the eighth grade of 10 percentage points (self-reported data) and over six years (Angrist et al, 2006), an increase in students completing secondary school of 15-20 percentage points (administrative data). Compounded over the relevant time-frame, these results are very close to our one-year enrollment effects for the savings and tertiary treatments assuming that the treatment effects do not decline over time, though unlike the PACES study, we are not yet able to assess the effects of continued enrollment on students' learning levels. While a more direct comparison will be possible as the students without cohort age, this similarity is intriguing because at a cost of US$43 per student per year, the savings and tertiary treatments are less than a quarter of the cost of the PACES vouchers (US$190 per student per year).

B. Other Outcomes

From students who completed both a baseline and follow-up survey, we 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 outcomes. 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, except for the rate at which students' pass. 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 individually. Again, however, when we pool the results, the overall average effect of the treatments is 1.5 percentage points which is statistically significant at the five percent level. Looking at the remaining outcome variables, students grades, none of the treatments cause a significant change.[14] Given the magnitude in the change in attendance rates, it seems reasonable that these changes may not reflect changes in students knowledge of material. However, it may be that teachers and administrators consider more than grades when determining who to promote to the next grade at the end of the academic year.

Panel B shows our findings for consumption and is divided into food and school expenses. The effects on food consumption (first two rows) are small but 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. 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 142,000 pesos anyway on average, but the treatment causes families to spend an extra 246,000 pesos. In total, it seems that 64.7 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

---

[14] It is interesting to note the relative similarity of the effects estimated through self-reported and verified grades. Despite the significant overestimates for enrollment and attendance, students reports of their grades are remarkably similar to those that we were able to verify. Self-reported grades are slightly higher, but measured treatment effects are very similar.

[18] 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.

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 percentage of kids whose primary activity is attending school by 12 percentage points, and reduces, by 4 percent points the fraction, 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 and earn 2,000 pesos less during the last week that they worked.

C. 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 7. We investigate 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, administrative enrollment, and the number of hours worked. 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 enrollment rates and columns 5 and 6 show the results for the number of hours worked. 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. 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 in Table 5, the treatment effect for students with a project attendance rate of less than 80 percent is slightly higher than that for students with a higher projected attendance rate. This pattern is consistent across all of the outcomes, although it is never statistically significant at conventional levels of significance.

The next two rows divide the sample by gender. In general, boys seem to respond more strongly than girls. Girls for example, respond to the treatments by attending 2.4 percent points more often while boys attend 3.6 percent more often.[18] Similarly, girls experience an increase in enrollment hat is 2.8 percentage points lower than that of boys, and the reduction in hours worked for girls is about 20 percent lower than that for boys.

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 in attendance rates (4.7 percentage points on the verified attendance measure) while those at the bottom show a much smaller response (1.4 percentage points of verified attendance). The same relative effects are observed for enrollment 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.

D. Peer Effects

Because our randomization strategy induces intra-school variation in the treatments, we are also able to measure directly the peer effects associated with the program 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 by looking for changes in the children who would not have been directly treated in both treatment and control villages. However, the relative percentage of treated and untreated children in each village (the treatment density) is determined by the distribution of wealth in each village, making it difficult to causally associate the magnitude of the effect with the density of treatment in a village.

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 faction of registered friends actually receiving the treatment. The question itself asked the children to name up to five of their closest friends who were in the same school and grade. All of the children provided at least 1 friend and three quarters provided five. The names of these friends were then matched with the results of the random assignment lottery to determine which friends were treated by the program. Half of the students had not treated friends and about 92 percent had up to 40 percent of their friends treated. This provides significant variation in treatment density, but also means that our results cannot be generalized to conditions in which more than 40 percent of a child's friends are treated. All of the results are presented in Table 8.

The first column of Table 8 contains a simple regression of the fraction of friends treated on the dummy variable for whether or not a student received a treatment directly or not. On average, 16 percent of a child's friends are treated (about 1 friend). This probability of a friend being treated is also equally divided between those students who were selected for treatment and those who were not. The coefficient on the treatment indicator variable is 0.6 percentage points and statistically insignificant despite the very small standard errors are on the coefficient (0.006).

The basic instrumental variables regression is laid out in columns 2, 3 and 4. Column 2 contains the OLS regression between our verified attendance measure and the fraction of a students' friends that are treated. Already the positive but declining contribution of treated friends is apparent but statistically insignificant. Column 3 contains the first stage regression that confirms that the outcome of the random assignment process (the fraction of registered friends who are treated) does, in fact, generate significant variance in the overall fraction of a student's friends who are treated. Finally Colum 4 contains the instrumental variables model.

The IV model shows that the OLS measure seems to be biased downwards as the instrumental variable regressions has larger and more statistically significant coefficients on the density of treatment within a child's friends. At the average treatment density, the overall magnitude of this effect is 1.2 percentage points, which is of the same magnitude of the direct effect. Combined with the direct effect, the total change in attendance, on average, in this sample of students is 2.3 percentage points. It is important to note, however, that while meaningful at the average treatment density, the indirect effects decline as additional friends are treated. So, there is a large gain to having a single friend treated (roughly the size of the average treatment density), but an additional friend contributes almost nothing.

In theory, peer effects may operate differently than the direct effect of the incentives. For example, if students are following a treated friend by attending more, this effect might not depend on the students' treatment status. We compare these relative effects in columns 5 and 6 which contain the results for students selected to receive a subsidy and those who were not. The results suggest that the indirect effect on both groups of students is similar. Although only those in the treatment regression are statistically significant at conventional levels (the join significance of the coefficients in the control regression are significant at the 13 percent level), the coefficients on the treatment density in both regressions are similar in magnitude.

In regressions not presented in this draft, we also estimated the effect of peer networks on enrollment using the samples in which we observed direct effects of the subsidies. In none of these regressions did we find a statistically significant effect of peers on the decision to enroll. However, because for students giving a follow-up survey we only observe direct effects of the subsidies in Suba and on students who would not have been likely to enroll without the subsidy, our inability to observe an effect may be due low statistical power as well as the fact that peers do not influence the decision to enroll in same way that they influence the decision to attend.

Finally, in column 7, we re-estimate the peer effects model using the verified attendance measure, but instead of the density of treatment within the peer group, we use the number of other students treated within a child's school-grade-gender cohort. While the direct effect is still precisely estimated, this measure of indirect effects does not detect the peer effects that we were able to measure with the explicit friendship mapping. The coefficients are very small in magnitude and statistically insignificant. These results suggest that indirect measure of peer effects may significantly underestimate the magnitude of such effects.

VI. Conclusion

This project demonstrates that both that the structure of educational incentives matter and that experiments involving these treatments 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. Overall the incentives make students directly treated by the program more likely to attend school (2.8 percentage points), more likely to remain enrolled (2.5 percentage points), more likely to matriculate to the next grade (1.6 percentage points), more likely to graduate (3.9 percentage points), and more likely to matriculate to a tertiary institution (21 percentage points). But within the treatments, however, the structure of the incentives matters. Simply changing the timing of the transfer with the savings incentive increases enrollment in both secondary and tertiary institutions without reducing daily attendance. The tertiary treatment is particularly effective, encouraging higher levels daily attendance (3.5 percentage points more for those students least likely to attend) and higher levels of enrollment at the secondary (3.3 percentage points) and tertiary levels (46 percentage points). These results suggest that savings constraints and the ability to enroll in tertiary institutions may be more important in the decision to keep students in school than short-term trade-offs within the academic year such as the opportunity cost of employment.

We also find that the direct incentive generate important externalities within peer networks. Treating friends encourages higher attendance at a similar magnitude as the direct effect. This effect appears to be independent of students' direct receipt of the incentive or of a students' counter-factual attendance rates. However, the gains in attendance decline sharply in the fraction of friends treated – one treated friend has significant benefits, but an additional friend has almost no additional 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.

Ashraf, Nava, Dean Karlan, and Wesley Yin (2006) "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Phillipines," Quarterly Journal of Economics, 121(2): 635-672.

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-Osorio, 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.

Blundell, Richard, Pierre-Andre Chiappori, and Costas Meghir (2005) "Collective Labor Supply With Children," *Journal of Political Economy*, 113(6): 1277-1306

Bobonis, Gustavo and Frederico Finan (2007). "Endogenous Peer Effects in School Participation" Manuscript. University of Toronto Department of Economics.

Bobonis, Gustavo, Edward Miguel, and Charu Sharma (2006) “Iron Deficiency Anemia and School Participation,” *Journal of Human Resources*.41(4), 692-721.

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 (2004). “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.

Kremer, Michael, Edward Miguel and Rebecca Thornton (2007) “Incentives to Learn,” Manuscript. University of California at Berkeley, Department of Economics.

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

Lalive R, and A. Cattaneo (2006). “Social Interactions and Schooling Decisions,” IZA Discussion Papers 2250, Institute for the Study of Labor (IZA).

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.

Miguel, Edward and Michael Kremer (2004) "Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities," *Econometrica*. 72(1): 159-217.

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) *Deevelopment 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 200,* The World Bank, Washington.

Figure 1: Distribution of Family Income

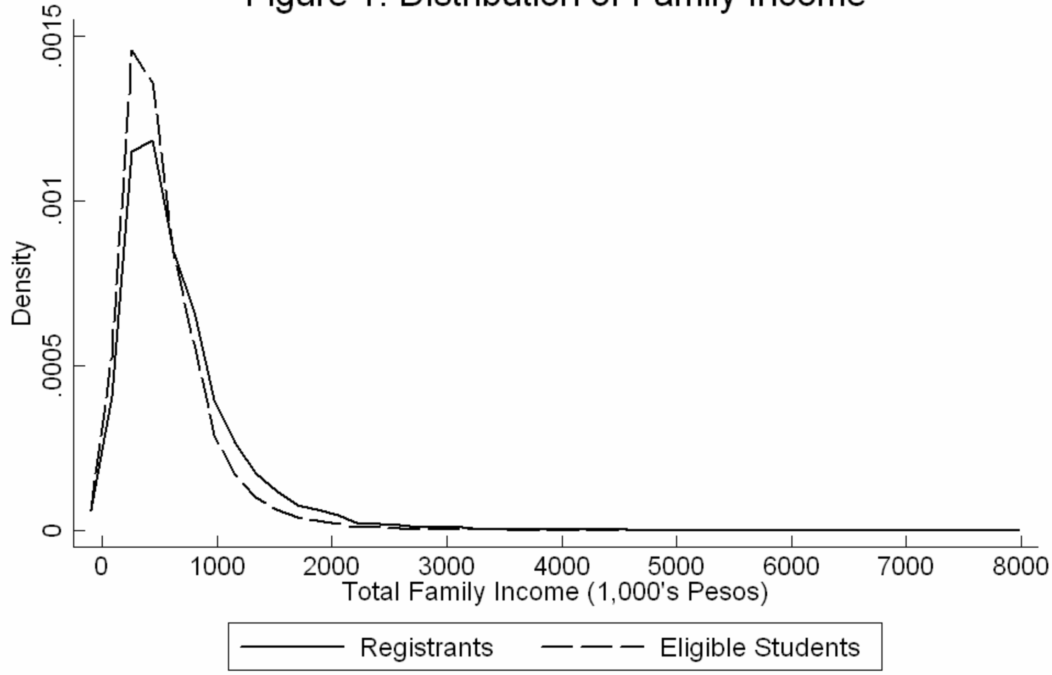

Figure 2: Distribution of Attendance

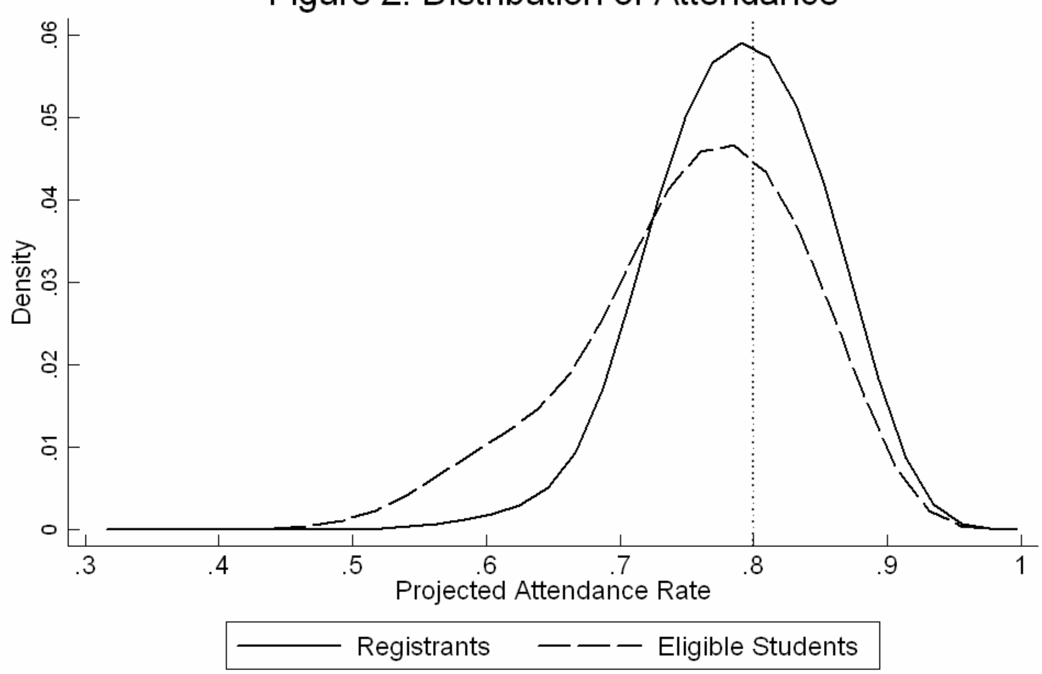

Figure 3: Distribution by Family Income at Baseline

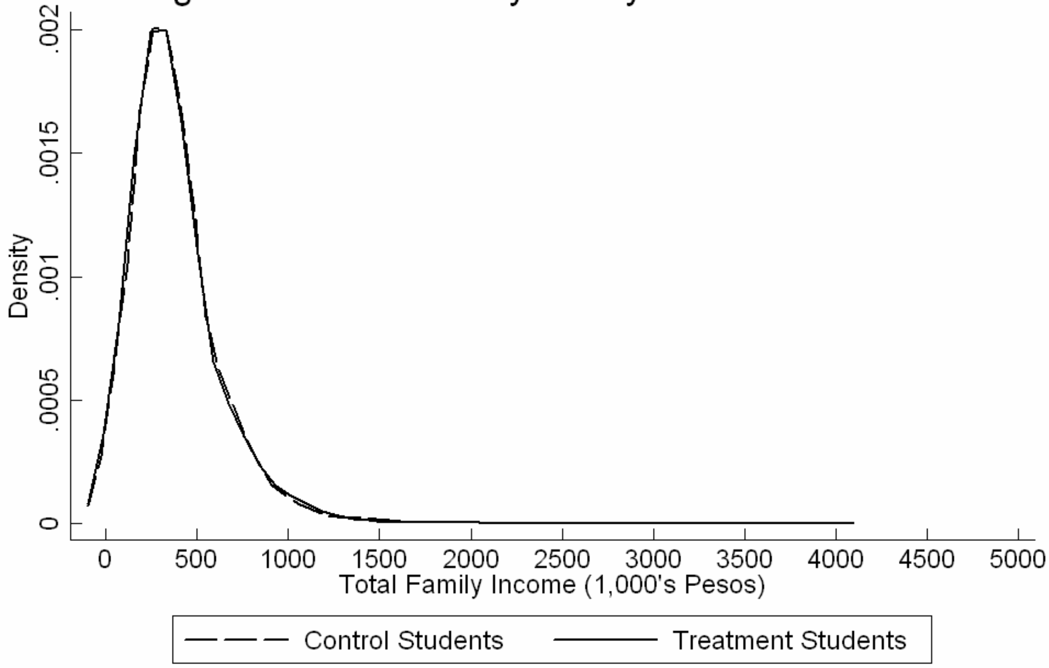

Figure 4: Distribution of Attendance at Follow-Up

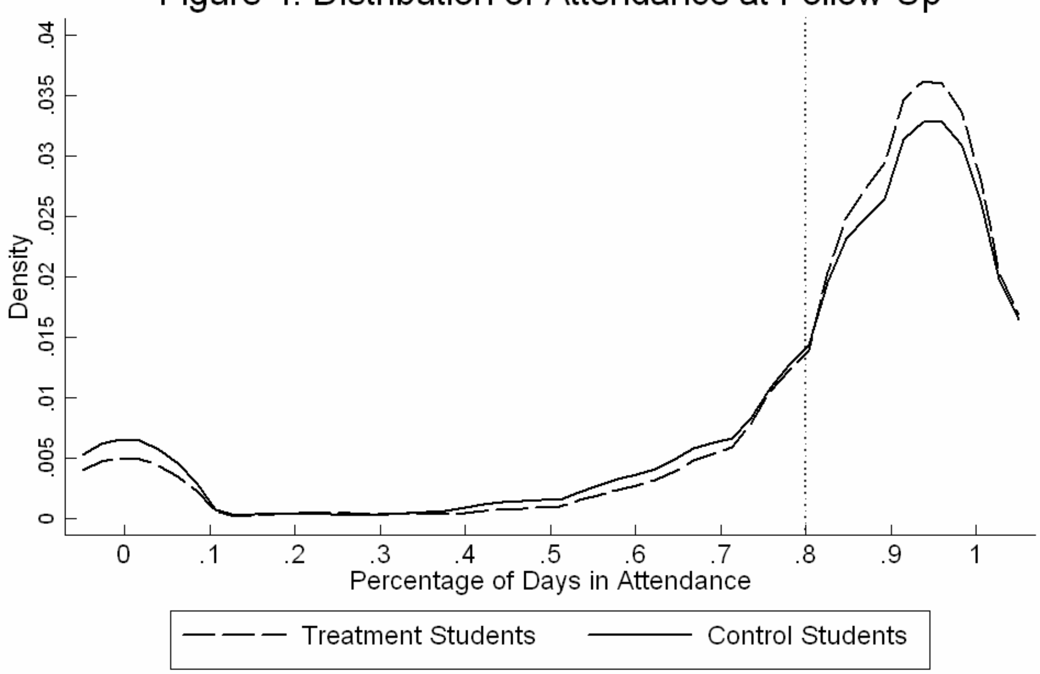

Figure 5: Actual Attendance by Predicted Attendance

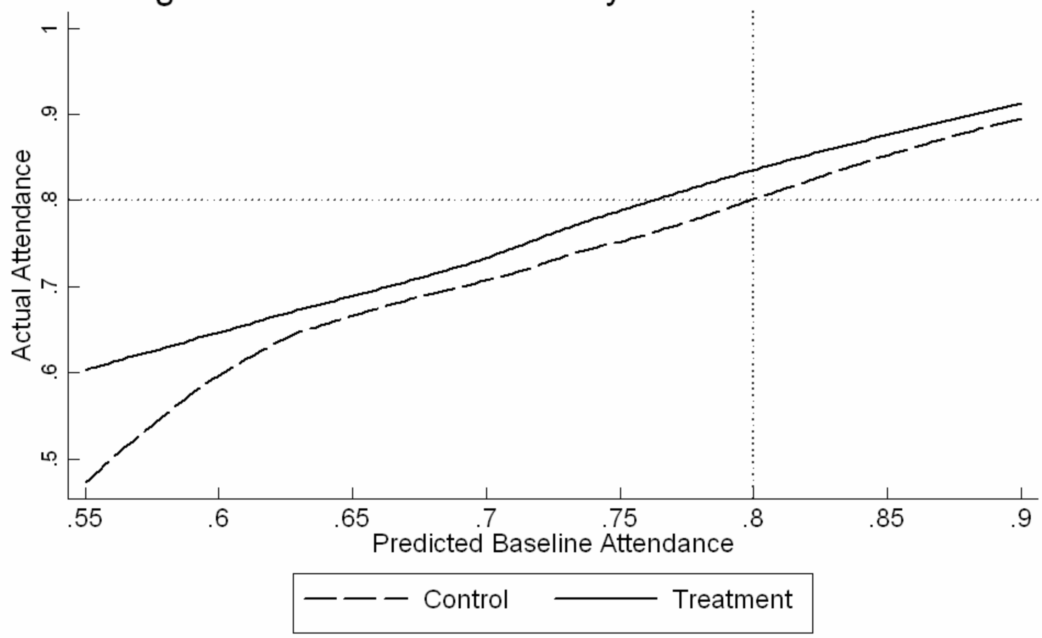

Note: Results from local polynomial regressions (bandwidth=0.075)

Figure 6: Enrollment by Predicted Enrollment

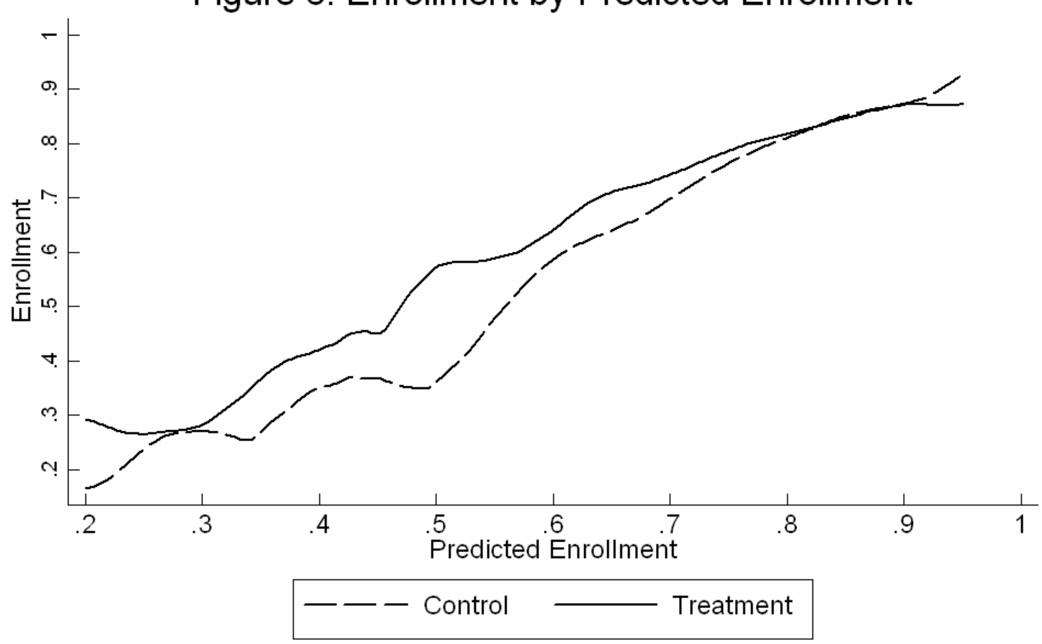

Note: Results from local polynomial regressions (bandwidth=0.075)

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	0.03**	0.08***	0.08***
	(1.09)	(0.01)	(1.00)	(0.01)	(0.01)	(0.02)
Utilities	4.77	-0.09***	4.85	-0.09***	-0.08***	-0.08***
	(1.40)	(0.01)	(1.35)	(0.02)	(0.01)	(0.02)
Durable Goods	1.5	-0.13***	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	-0.24***	-0.34***	-0.33***
	(1.79)	(0.02)	(1.58)	(0.02)	(0.01)	(0.03)
Panel B: Individual Characteristics						
Age	15.16	-1.53***	15.08	-1.49***	-0.01	0.04
	(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	0.15***	0.01**	0
	(0.77)	(0.01)	(0.77)	(0.01)	(0.00)	(0.00)
Attending School	75.61	20.08***	76.15	18.93***	0.80***	0.61*
	(42.95)	(0.43)	(42.62)	(0.55)	(0.24)	(0.33)
Years of Education	6.33	-0.71***	6.5	-0.69***	-0.22***	-0.19***
	(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)	(0.01)	(0.46)	(0.01)	(0.00)	(0.01)
Age of Head	45.97	0.01	44.77	0.43***	1.15***	0.78***
	(11.05)	(0.12)	(10.03)	(0.13)	(0.06)	(0.16)
Years of Ed, Head	5.91	-0.35***	6.33	-0.46***	-0.43***	-0.30***
	(3.09)	(0.03)	(3.16)	(0.04)	(0.02)	(0.05)
People in Household	5.27	0.13***	5.07	0.12***	0.22***	0.23***
	(2.06)	(0.02)	(1.89)	(0.03)	(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	-1.16***	14.28	-1.10***	-1.47***	-1.45***
	(5.01)	(0.05)	(4.74)	(0.06)	(0.03)	(0.07)
Household Income	437.04	-71.19***	482.96	-84.85***	-48.39***	-32.26***
(1,000 Pesos)	(282.34)	(2.95)	(295.62)	(3.91)	(1.66)	(3.82)

* 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

	San Cristobal				Suba (Grade 6-8)		Suba (Grade 9-11)	
Demographic Variable	**Control Average**	**Basic-Control**	**Savings-Control**	**Basic-Savings**	**Control Average**	**Basic-Control**	**Control Average**	**Tertiary-Control**
Panel A: Indexes of Household Assets								
Posessions	1.94	0.06*	0	0.06*	1.84	-0.02	1.93	0.03
	(1.11)	(0.03)	(0.03)	(0.04)	(1.00)	(0.05)	(1.01)	(0.06)
Utilities	4.67	-0.03	0.04	-0.06	4.68	0.06	4.86	0.08
	(1.42)	(0.05)	(0.05)	(0.04)	(1.40)	(0.06)	(1.31)	(0.06)
Durable Goods	1.37	-0.03	0.02	-0.05**	1.54	0.01	1.62	0.07
	(0.88)	(0.03)	(0.03)	(0.02)	(0.86)	(0.04)	(0.84)	(0.05)
Physical Infrastructure	11.64	-0.10**	0.01	-0.11*	11.92	0	12.12	0.01
	(1.72)	(0.05)	(0.04)	(0.06)	(1.45)	(0.04)	(1.38)	(0.08)
Panel B: Individual Characteristics								
Age	14.36	0.01	-0.22	0.23	12.65	0.05	15.54	0.25
	(5.51)	(0.12)	(0.21)	(0.19)	(3.89)	(0.16)	(4.18)	(0.31)
Gender	0.49	0.02	-0.01	0.03**	0.48	0.01	0.42	0
	(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	4.96	-0.02
	(0.53)	(0.01)	(0.02)	(0.02)	(0.30)	(0.01)	(0.38)	(0.03)
Years of Education	5.6	-0.10**	-0.06	-0.04	4.66	0.04	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	0	0.01	0.27	0	0.25	0.03
	(0.46)	(0.01)	(0.01)	(0.02)	(0.44)	(0.01)	(0.43)	(0.02)
Age of Head	45.76	-0.1	0.09	-0.18	44.69	-0.22	45.75	0.79**
	(10.33)	(0.27)	(0.36)	(0.30)	(9.11)	(0.41)	(8.54)	(0.32)
Years of Ed, Head	5.59	-0.07	-0.07	0	5.69	0.03	5.79	-0.1
	(2.90)	(0.10)	(0.08)	(0.10)	(2.79)	(0.09)	(2.97)	(0.13)
People in Household	5.39	0.03	0.02	0.01	5.26	-0.02	5.13	0.1
	(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	2.36	0.03
	(1.32)	(0.04)	(0.04)	(0.04)	(1.29)	(0.06)	(1.18)	(0.09)
Panel D: Poverty Measures								
Estrato	1.45	-0.01	0	-0.01	1.57	0	1.65	0
	(0.82)	(0.03)	(0.03)	(0.02)	(0.82)	(0.03)	(0.75)	(0.04)
SISBEN Score	11.75	-0.24*	-0.15	-0.09	13.11	-0.03	13.51	0.24
	(4.63)	(0.14)	(0.15)	(0.13)	(4.37)	(0.12)	(4.24)	(0.27)
Household Income	364.27	-2.48	5.26	-7.74	390.2	1.19	396.13	4.18
(1,000 Pesos)	(235.66)	(7.49)	(8.66)	(7.07)	(223.22)	(9.62)	(228.19)	(12.43)

* 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

	San Cristobal				Suba (Grade 6-8)		Suba (Grade 9-11)	
Demographic Variable	**Control Average**	**Basic-Control**	**Savings-Control**	**Basic-Savings**	**Control Average**	**Basic-Control**	**Control Average**	**Tertiary-Control**
Panel A: Attrition Rate								
Number Attritors	44	-3	-17		18	0	8	1
Percentage of Baseline	0.02	0	-0.01	0.01	0.02	0	0.01	0
	(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*	-0.1	-0.17	0.07	-0.49**	-0.01	-0.31	0.36
	(0.17)	(0.17)	(0.23)	(0.17)	(0.24)	(0.19)	(0.36)	(0.47)
Utilities	-0.75***	0.43	0.38	0.05	-0.63*	-0.54	0.14	-1.44***
	(0.22)	(0.30)	(0.25)	(0.30)	(0.34)	(0.42)	(0.47)	(0.51)
Durable Goods	-0.21	-0.01	-0.11	0.1	-0.55***	0.18	-0.50*	-0.08
	(0.13)	(0.22)	(0.21)	(0.21)	(0.21)	(0.25)	(0.30)	(0.33)
Physical Infrastructure	-0.31	-0.25	-0.45	0.2	-0.16	-0.44	0.13	-1.61**
	(0.26)	(0.37)	(0.43)	(0.47)	(0.35)	(0.48)	(0.49)	(0.70)
Panel C: Individual Characteristics								
Age	2.63***	-2.39	-2.87*	0.48	-0.13	0.51	0.47	-1.15*
	(0.84)	(2.11)	(1.60)	(1.29)	(0.95)	(0.44)	(1.49)	(0.64)
Gender	0.11	-0.15	-0.13	-0.02	0.23*	-0.27***	-0.30*	0.1
	(0.08)	(0.11)	(0.13)	(0.12)	(0.12)	(0.09)	(0.18)	(0.20)
Married	-0.06	0.03	0.12	-0.09	0.02	0.02	0.04	0.01
	(0.08)	(0.15)	(0.09)	(0.10)	(0.07)	(0.02)	(0.14)	(0.03)
Years of Education	-0.28	0.32	0.28	0.04	0.04	0.53	1.32***	-1.25**
	(0.29)	(0.45)	(0.68)	(0.53)	(0.31)	(0.39)	(0.44)	(0.57)
Panel D: Household Characteristics								
Single Head	-0.09	0.19**	0.19	0	0.09	0.15	0.12	0.03
	(0.07)	(0.09)	(0.12)	(0.15)	(0.11)	(0.13)	(0.15)	(0.17)
Age of Head	-1.45	-3.77**	2.97	-6.74***	-3.63	2.07	0	-6.18*
	(1.57)	(1.75)	(2.27)	(2.52)	(2.23)	(2.21)	(3.04)	(3.38)
Years of Ed, Head	0.89**	-0.82	-0.81	-0.01	0.55	-1.72*	0.6	-0.57
	(0.44)	(0.64)	(0.65)	(0.82)	(0.68)	(0.98)	(1.06)	(1.28)
People in Household	0.06	-0.4	-0.21	-0.19	-0.26	-0.65	-0.38	-0.41
	(0.29)	(0.34)	(0.38)	(0.33)	(0.43)	(0.52)	(0.58)	(0.81)
Member under 18	-0.09	0.29	0.24	0.05	0.25	-0.48	0.14	-0.34
	(0.20)	(0.38)	(0.39)	(0.25)	(0.32)	(0.38)	(0.42)	(0.75)
Panel E: Poverty Measures								
Estrato	-0.34***	0.17	0.24	-0.07	-0.04	-0.5	0.1	-0.99***
	(0.12)	(0.19)	(0.22)	(0.20)	(0.20)	(0.31)	(0.27)	(0.31)
SISBEN Score	-1.56**	0.97	1.01	-0.04	-1.09	-1.81	-0.82	-2.85*
	(0.71)	(0.99)	(1.06)	(1.31)	(1.07)	(1.36)	(1.51)	(1.52)
Household Income	-0.86	-46.84	-35.24	-11.61	-104.63**	82.87	-87.77	78.28
(1,000 Pesos)	(35.93)	(49.74)	(55.87)	(46.03)	(52.99)	(66.61)	(81.20)	(65.00)

* 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

	San Cristobal			Suba, Grades 6-8		Suba, Grades 9-11		All
Outcome Variable	**Control Average**	**Basic - Control**	**Savings - Control**	**Control Average**	**Basic - Control**	**Control Average**	**Tertiary - Control**	**Treat-Control**
Panel A: Verified Attendance, 2005								
All	0.794	0.033***	0.028***	0.782	0.009	0.793	0.050***	0.028***
	(0.006)	(0.006)	(0.006)	(0.009)	(0.012)	(0.011)	(0.015)	(0.005)
Grades 6-8	0.792	0.035***	0.026***	0.782	0.009			0.023***
	(0.008)	(0.009)	(0.009)	(0.009)	(0.012)			(0.007)
Grades 9-11	0.797	0.030***	0.030**			0.793	0.050***	0.036***
	(0.009)	(0.010)	(0.012)			(0.011)	(0.015)	(0.008)
Baseline Att >= 0.8	0.857	0.024**	0.022***	0.878	0.021*	0.86	0.032**	0.024***
	(0.007)	(0.010)	(0.007)	(0.009)	(0.012)	(0.012)	(0.014)	(0.006)
Baseline Att < 0.8	0.728	0.043***	0.033***	0.63	-0.009	0.688	0.078**	0.035***
	(0.009)	(0.014)	(0.011)	(0.016)	(0.017)	(0.019)	(0.034)	(0.010)
Panel B: Verified Attendance if Followup, 2005								
All	0.872	0.011**	0.013**	0.841	0.007	0.857	0.016	0.011**
	(0.004)	(0.005)	(0.006)	(0.007)	(0.008)	(0.008)	(0.013)	(0.004)
Grades 6-8	0.874	0.012*	0.011*	0.841	0.007			
	(0.005)	(0.007)	(0.006)	(0.007)	(0.008)			
Grades 9-11	0.869	0.011	0.016*			0.857	0.016	0.014*
	(0.006)	(0.008)	(0.010)			(0.008)	(0.013)	(0.007)
Baseline Att >= 0.8	0.912	0.005	0.005	0.921	0.015	0.91	0.009	0.007**
	(0.004)	(0.004)	(0.006)	(0.005)	(0.009)	(0.007)	(0.011)	(0.004)
Baseline Att < 0.8	0.825	0.018*	0.022**	0.7	-0.001	0.766	0.028	0.016**
	(0.006)	(0.010)	(0.010)	(0.014)	(0.011)	(0.015)	(0.026)	(0.008)
Panel C: Enrollment, 2006								
All Students	0.698	0.009	0.036***	0.704	0.027*	0.741	0.033*	0.025***
	(0.008)	(0.010)	(0.011)	(0.010)	(0.014)	(0.014)	(0.019)	(0.006)
All, Pred Enroll >= 0.8	0.847	-0.005	-0.008	0.853	-0.021	0.886	-0.009	-0.009
	(0.012)	(0.020)	(0.017)	(0.014)	(0.019)	(0.015)	(0.021)	(0.011)
All, Pred Enroll < 0.8	0.639	0.016	0.056***	0.63	0.049***	0.61	0.080***	0.042***
	(0.010)	(0.014)	(0.011)	(0.014)	(0.017)	(0.022)	(0.027)	(0.008)
Completed Survey	0.834	-0.002	0.008	0.827	0.024	0.857	0.024	0.01
	(0.009)	(0.012)	(0.014)	(0.012)	(0.016)	(0.017)	(0.017)	(0.009)
Self Reported	0.982	0	0.001	0.994	-0.003	0.988	0.008	0.001
	(0.003)	(0.003)	(0.004)	(0.003)	(0.004)	(0.005)	(0.005)	(0.002)
Panel C: Self Reported Attendance, 2006								
All	0.958	0.005	0.006	0.962	0.010**	0.955	0.016*	0.008**
	(0.003)	(0.004)	(0.004)	(0.003)	(0.004)	(0.005)	(0.009)	(0.003)
Baseline Att >= 0.8	0.964	0.004	0.004	0.964	0.007	0.968	0.01	0.005*
	(0.003)	(0.004)	(0.004)	(0.003)	(0.005)	(0.005)	(0.006)	(0.003)
Baseline Att < 0.8	0.95	0.005	0.008	0.957	0.014**	0.932	0.022	0.010**
	(0.004)	(0.006)	(0.005)	(0.005)	(0.005)	(0.011)	(0.020)	(0.005)
Panel E: Students in Final Year of Secondary School in 2005								
Graduated, 2005	0.876	0.029	0.043			0.903	0.054	0.039*
	(0.025)	(0.044)	(0.032)			(0.024)	(0.040)	(0.023)
Higher Ed, 2006	0.227	0.04	0.088***			0.193	0.497***	0.209***
	(0.032)	(0.034)	(0.033)			(0.032)	(0.044)	(0.050)

* 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

	San Cristobal			Suba, Grades 6-8		Suba, Grades 9-11		All
Outcome Variable	**Control Average**	**Basic - Control**	**Savings - Control**	**Control Average**	**Basic - Control**	**Control Average**	**Tertiary - Control**	**Treat-Control**
Panel A: Academic Effort, Grades 6-11								
Hours of Homework	2.697	0.02	0.035	2.961	0.022	2.609	0.540***	0.098***
	(0.033)	(0.043)	(0.048)	(0.040)	(0.044)	(0.072)	(0.110)	(0.036)
Total Grades, Self Reported	0	0.060*	0.05	0.024	-0.04	-0.041	-0.046	0.021
	(0.024)	(0.034)	(0.036)	(0.034)	(0.062)	(0.040)	(0.052)	(0.025)
Total Grades, Verified	0	0.083	0.048	0.049	0.02	-0.097	-0.059	0.037
	(0.034)	(0.054)	(0.046)	(0.048)	(0.061)	(0.056)	(0.081)	(0.033)
Passed in 2005	0.889	0.01	0.019**	0.906	0.015	0.903	0.022	0.016**
	(0.007)	(0.009)	(0.010)	(0.009)	(0.014)	(0.012)	(0.017)	(0.006)
Panel B: Consumption, Grades 6-11								
Meals Over Last 3 Days	8.018	0.191**	0.239***	8.197	0.065	8.199	0.166**	0.177***
	(0.051)	(0.076)	(0.073)	(0.059)	(0.092)	(0.075)	(0.076)	(0.049)
Meals with Eggs or Meat	5.069	0.164**	0.182***	5.24	0.05	5.278	0.172*	0.144***
	(0.042)	(0.074)	(0.053)	(0.050)	(0.070)	(0.063)	(0.092)	(0.043)
School Expenses	231.419	9.997	20.972***	236.285	-13.348	269.817	-20.706	4.228
Grades 6-10	(5.090)	(7.105)	(6.311)	(6.815)	(10.182)	(11.440)	(12.656)	(5.471)
School Expenses	171.559	8.584	19.878			142.329	246.038***	95.729***
Grade 11	(31.640)	(37.512)	(44.919)			(30.014)	(47.523)	(33.638)
Panel C: Labor Activities, Grades 6-11								
Primary Activity, Studying†	0.879	0.001	0.006	0.913	-0.002	0.734	0.127***	0.018**
	(0.007)	(0.007)	(0.008)	(0.009)	(0.013)	(0.018)	(0.025)	(0.008)
Primary Activity, Work	0.021	-0.002	0.001	0.005	0.002	0.071	-0.041***	-0.005**
	(0.003)	(0.003)	(0.005)	(0.002)	(0.003)	(0.010)	(0.009)	(0.003)
Primary Activity, Home	0.022	0.003	0.005	0.005	0.002	0.056	-0.041***	-0.002
	(0.003)	(0.004)	(0.004)	(0.002)	(0.004)	(0.009)	(0.009)	(0.003)
Hours Worked	1.459	-0.335*	-0.219	0.626	-0.178	3.358	-2.067***	-0.492***
Last Work Week	(0.169)	(0.202)	(0.212)	(0.117)	(0.138)	(0.459)	(0.548)	(0.141)
Earnings	2.032	-0.246	-0.067	0.793	0.17	5.708	-2.751**	-0.446
Last Work Week	(0.264)	(0.394)	(0.400)	(0.146)	(0.441)	(0.932)	(1.238)	(0.323)

†Percentages for each primary activity do not add to 100 percent because two categories are omitted (incapacitated, and other activities). Standard errors are clustered at the family level. * significant at 10 percent level, ** at 5 percent level, and *** at 1 percent level.

Table 7: Heterogeneity of Treatment Effects

Demographic Variable	Verified Attendance: Control Average	Verified Attendance: Treatment - Control	Admin Enrollment: Control Average	Admin Enrollment: Treatment - Control	Hours Worked (Week): Control Average	Hours Worked (Week): Treatment - Control
All Students	0.791	0.029***	0.709	0.023***	1.589	-0.479***
	(0.015)	(0.005)	(0.050)	(0.006)	(0.151)	(0.147)
Baseline Attendance						
Attendance > 0.8	0.862	0.024***	0.703	0.016***	1.421	-0.432***
	(0.006)	(0.006)	(0.067)	(0.006)	(0.214)	(0.159)
Attendance ≤ 0.8	0.705	0.036***	0.726	0.044***	1.809	-0.524**
	(0.018)	(0.010)	(0.018)	(0.015)	(0.265)	(0.245)
Gender						
Female	0.803	0.016**	0.73	0.015*	1.205	-0.345**
	(0.017)	(0.006)	(0.052)	(0.009)	(0.180)	(0.162)
Male	0.778	0.043***	0.687	0.034***	2.008	-0.588**
	(0.015)	(0.008)	(0.050)	(0.009)	(0.189)	(0.260)
Income						
Upper Tercile	0.779	0.047***	0.705	0.039***	1.521	-0.619*
	(0.019)	(0.010)	(0.054)	(0.011)	(0.280)	(0.335)
Middle Tercile	0.801	0.023***	0.716	0.025*	1.594	-0.666***
	(0.014)	(0.008)	(0.046)	(0.013)	(0.236)	(0.206)
Low Tercile	0.794	0.017*	0.705	0.01	1.652	-0.38
	(0.015)	(0.009)	(0.054)	(0.012)	(0.240)	(0.247)

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

Table 8: Effects of Peer Networks on Verified Attendance Measure

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dependent Variable	**Fraction of Treated Friends**	**Verified Attendance**	**Fraction of Treated Friends**	**Verified Attendance**	**Verified Attendance**	**Verified Attendance**	**Verified Attendance**
Model	**OLS**	**OLS**	**OLS**	**IV**	**IV**	**IV**	**IV**
Observations Used	All	All	All	All	Treatment	Control	All
Treatment	0.0062	0.0110**	0.0007	0.0111**			0.0107**
	(0.006)	(0.004)	(0.003)	(0.004)			(0.004)
Fraction of Friends Treated		0.0252		0.1451**	0.1695***	0.1309	
		(0.020)		(0.061)	(0.057)	(0.112)	
Fraction of Friends Treated		-0.0318		-0.3291**	-0.4196***	-0.2636	
Squared		(0.032)		(0.147)	(0.137)	(0.267)	
Fraction of Registered Friends			0.7643***				
Treated			(0.015)				
Fraction of Registered Friends			-0.4271***				
Treated Squared			(0.017)				
Number of Students Treated							-0.0006
by Grade-Gender Cohort							(0.010)
Number of Treated Squared							0
by Grade-Gender Cohort							(0.001)
Student-Family Controls		√	√	√	√	√	√
School Fixed Effects		√	√	√	√	√	√
Constant	0.1574***	0.0403	0.1024	0.0614	-0.0222	1.1612***	0.0478
	(0.009)	(0.102)	(0.072)	(0.101)	(0.126)	(0.152)	(0.129)
R-squared	0	0.25	0.69	0.24	0.25	0.24	0.25

†Treatment density for friendship network measure is the percentage of friends treated. For Grade-Gender Groups, treatment desnity is the number of students treated. Standard errors are clustered at the school level. * significant at 10 percent level, ** at 5 percent level, and *** at 1 percent level