Child Control in Education Decisions: An Evaluation of Targeted Incentives to Learn in India

James Berry[*,†]
Cornell University

September 2009

Abstract

In developing countries, a common approach to encourage education is to provide cash incentives to households conditional on their children's enrollment, attendance or achievement in school. Little is known, however, about whether the identity of recipient—either the parent or child—can influence the effectiveness of these interventions. This study provides an empirical test of whether changing the recipient of the incentive affects outcomes. To structure the test, I develop a model of household education production in which parents' ability to motivate their children is dampened because of moral hazard. In the model, when parental inputs are more productive, rewarding parents is more effective than rewarding children, while the opposite is true when child inputs are more productive. A testable prediction is that the relative effectiveness of rewarding the parent rather than the child should be positively correlated with the part of the child's prior achievement that is predicted by parental characteristics. I test the model with a field experiment in Gurgaon, India where an incentive to achieve a specific reading goal was randomly assigned to be received by either the parent or by the child. I find that incentives to parents result in worse outcomes than incentives to children when children have less productive parents and lower initial test scores. Conversely, incentives to parents result in better outcomes when children have more productive parents and higher initial test scores.

JEL classification: C93, D13, I21

Keywords: student incentives, household decision making, education policy, field experiments

[*] I am especially grateful to my advisors, Esther Duflo, Abhijit Banerjee and Ben Olken as well as Nava Ashraf, Neil Bhutta, Camilo Dominguez, Greg Fischer, Raymond Guiteras, Cynthia Kinnan, Jeanne Lafortune, Leigh Linden, Karthik Muralidharam, Trang Nguyen, Jesse Shapiro, Christopher Smith, David Tabak, José Tessada and Tom Wilkening for helpful comments and suggestions. I would also like to thank Hansa Parmar for her tireless work coordinating the field activities in Gurgaon, Swati Gaur for research assistance, the staff at Pratham Gurgaon, especially Brij Kaul, Pawan Sharma and Rajeev Ranjan, and the Gurgaon district officer for his support. All errors are my own. This project was supported by the Russell Sage Foundation Small Grants Program in Behavioral Economics and the George and Obie Schultz Fund.

[†] E-mail: *jwb295@cornell.edu*

1 Introduction

In developing countries, increasing enrollment and learning in primary school remains an important policy goal. While many countries have made progress towards universal primary enrollment, 14% of primary-aged children in the developing world are still out of school (UNESCO, 2007). In addition, these countries still face the challenge of educating the children who are in school. A recent survey in India, for example, found that only 50% of children in grade 3 could read at a first-grade level (Pratham, 2008).

Interventions to improve schooling outcomes often take the form of supply-side reforms, such as improvements in infrastructure, materials or teachers. Recently an increasing amount of attention has been given to demand-side interventions, which either lower the costs or increase the immediate benefits that households face when deciding to educate their children. One common demand-side intervention offers cash rewards to households when their children enroll in, attend or achieve in school.

While interventions that provide cash rewards condition the transfer on different behaviors (enrollment, attendance or achievement), those that target families in developing countries primarily provide the rewards to parents. Implicit in the design of these interventions is the idea that rewarding the parents, rather than the children, produces the best results. In contrast to a literature suggesting that the recipient of a transfer matters between adults within a household (Duflo, 2003; Lundberg et. al., 1997), to my knowledge there have been no comparisons of programs that target parents with those that target children.

Changing the recipient of the incentive from the parent to the child may affect outcomes if child effort is an important input into education production and if parents are unable to motivate their children to the full extent they are themselves motivated by the incentives. If this is the case, incentives provided to children may result in better outcomes than incentives provided to parents. As I explore in this paper, the transfers that parents provide their children may be dampened either by agency problems within the household or by parents' inability to commit to rewarding their children for positive outcomes.

To my knowledge, this study is the first to test empirically whether providing an incentive to the child or to the parent has an effect on educational outcomes. In this paper I present the results of a field experiment in India that offered incentives targeted to either the child or the parent. By offering incentives directly to the child, the experiment made the child the full residual claimant of his[1] efforts. I compare the resulting education outcomes with the case when an equally-valued incentive was provided to the parent.

To motivate the experimental design, I develop a model of education production in which both the parent and child exert costly effort, and households vary in the productivity of the parent's and child's inputs in the production function. In the model, the productivity of parent and child inputs determines both the child's initial learning level and the relative effectiveness of incentives given to the parent versus incentives given to the child. The model has different implications for the relationship between initial test scores and relative effectiveness of incentives given to the parent, depending on whether the primary source of variation in initial test scores is parental or child productivity. If differences in initial test scores are driven by differences in parental productivity, there should be a positive association between the relative effectiveness of incentives given to the parent and initial test scores. In particular, children with low-productivity parents will be initially low achieving, and these children will perform better when they receive incentives directly compared to the case when their parents receive the incentives. Conversely, children with high-productivity parents will be initially high achieving, and these children will perform better when their parents receive the incentives. If initial test scores are instead driven by differences in child productivity, there should be a negative association between the relative effectiveness of incentives given to the parent and initial test scores.

I also present an extension to the model in which parents cannot commit to rewarding their children for learning. If parents cannot commit, children will not exert effort. The model implies that offering parents a commitment device that provides external incentives to their children will improve outcomes when the child's input is relatively more productive. Depending on whether variation in test scores is driven by parent or child productivity, parents of initially low- or high-performing children will be more willing to commit.

I test the model using a field experiment with primary school students in urban slums in Gurgaon, India. The Indian context is well suited to the purposes of this study. While net primary

[1] Throughout the paper, I use masculine pronouns to refer to the child and feminine pronouns to refer to the parent.

school enrollment improved from 79% in 2000 to 88% in 2005 (World Bank, 2008), attendance and performance in school still lags. Among enrolled children, absence rates of 25-35% are common (see, e.g., Banerjee et. al, 2005; Pratham, 2008). As noted above, many enrolled students also suffer from poor learning outcomes.

The available data suggest that parent and child effort, not financial considerations such as direct costs of school or opportunity costs of not working, may be an important factor in the lack of attendance and performance in Indian primary schools. There are no school fees, and some families receive subsidies for their children's enrollment. Children also receive free lunch at school through the Mid-day Meal program. Of the few children who are not in school, the majority do not work: according to the 2001 census, only 5% of children between the ages of 5 and 14 were working.[2]

The experiment offered prizes to the parent or child if the child reached a literacy goal after two months. Each child was given a goal based on his pretest score and was tested again after two months to determine if the goal had been reached. In order to isolate the effects of changing control over the rewards between the parent and child, program families were randomly offered either money given to the child's mother ("parent money"), money to the child ("child money"), or toys to the child as a reward. The parent money treatment gave the parent full control over the reward. In the child money treatment, the child was physically handed the reward, but the money could easily be taken from the child to be spent at the parent's discretion. In the toys treatment, the child was handed the reward, *and* he was given a good that the parent could not use herself. Because the parent could not easily capture the toy, and because baseline survey responses indicated that money given to the child was simply transferred to the parent, I focus on the difference between the aggregated money treatments and toys treatment in the analysis. Since the interaction of pretest score and the effect of parent incentives relative to child incentives is a key interaction of interest, the randomization was stratified by pretest score, ensuring balance of the treatment groups within each test score level.

The experiment included two additional treatments to test explicitly whether parents want to

[2] Among children enrolled in school, there also does not appear to be any tension between school and work. In the sample of in-school children used in this study, less than 1% reported participating in non-domestic work more than one hour per day, even though the school day ends at 12:30pm.

reward their children for performing well but cannot commit to doing so. I offered the parents a choice of money for themselves or a toy for their child, either *ex ante* (at the outset of the program) or *ex post* (after the goal had been reached). The *ex ante* treatment gave the parents the opportunity to commit, and the *ex post* treatment served to check that the choice itself was not driving outcomes.

Children in the program were given the opportunity to attend free after-school classes to assist them in improving their reading skills. These classes were held to give the children a greater chance to achieve the goals set out by the program. In addition, attendance in these classes provides an objective measure of effort and serves as an intermediate outcome in the analysis.

While my design did not allow for a pure control group, I present quasi-experimental evidence that the incentives program had a substantial impact on test scores across all incentive treatments. Turning to the differences between treatment groups, I find no significant mean differences in attendance or achievement between the three main treatment groups of parent money, child money and toys. I do, however, find substantial heterogeneity in treatment effects consistent with differences in parental productivity in education production. Children with lower initial test scores perform better when provided a toy as an incentive relative to money, while the reverse is true for children with higher initial test scores. To isolate the variation in test scores attributable to parental productivity, I construct an index of parental productivity using the predicted values from a regression of initial test scores on parental characteristics. I find that the effectiveness of the toys treatment is more negatively correlated with the portion of test scores related to parental productivity than with test scores overall. These results hold for attendance in the supplemental classes as well as for achievement of the literacy goal. Finally, I find no evidence that offering parents the opportunity to commit to rewarding their children with toys improves outcomes. This suggests that poor learning outcomes do not result from an inability to commit on the part of the parents.

This paper offers several contributions to existing research. As an study of an incentives-to-learn program, it adds to a new but growing literature evaluating these types of interventions. To my knowledge, the only randomized evaluation of such a program in developing countries is Kremer, Miguel and Thornton (2004). The authors evaluate a cash incentive program for primary-aged girls in Kenya. The program was structured as a tournament, where prizes were awarded to parents of girls who scored in the top 15% of a standardized test. The authors find a significant impact of the program on learning outcomes.

Several studies in developed countries, involving children at upper-secondary or tertiary levels of education, evaluate programs that rewarded students directly. In these studies the evidence on the effectiveness of the incentives is somewhat mixed. Angrist and Lavy (2002) find a significant impact of one of two evaluations of rewards for performance on high-school matriculation exams in Israel. Angrist, Lang and Oreopoulos (2006) find that a program rewarding good grades in a Canadian university raised grades for females, but not for males. In a recent study, Jackson (2008) finds that incentives for performance on Advanced Placement exams in the United States had a significant impact on performance on these tests.[3]

This paper also relates to the literature evaluating conditional cash transfer programs that provide families in developing countries with cash if the children enroll in or attend school. These evaluations generally find that conditional cash transfers achieve the objectives of increasing enrollment and attendance (Schultz, 2004; Glewwe and Olinto, 2004; Schady and Araujo, 2006; and Barrera-Osorio et. al., 2008).

More generally, this paper contributes to the theoretical and empirical literature on how parents provide incentives to their children. Becker's Rotten Kid Theorem (1974) provides an early theoretical backbone to this line of research. The Theorem shows that under certain assumptions, a parent can control her child's actions indirectly through transfers to her child. Through these transfers, the parent can make the child fully internalize the parent's value of schooling. External incentives provided to the parent will therefore produce equivalent results to incentives provided to the child. However, it has been noted that the conditions under which the Theorem holds are somewhat restrictive. Bergstrom (1989) points out a number of situations under which the theorem does not apply, including moral hazard. The model I present in the next section includes this feature. Gatti (2005) explores the theoretical implications of a moral hazard model for bequests and intergenerational transfers between parents and children. Weinberg (2001) finds some empirical evidence in favor of a moral hazard model by examining the relationship between household

[3] Note, however, that the Advanced Placement Incentive Program bundled both teacher and student incentives, and it is unclear how each type of incentive contributed to the success of the students.

income and the use of corporal punishment.

The remainder of this paper is organized as follows. Section 2 lays out a model of education production that motivates the experiment's design. Section 3 describes the design of the intervention and outcome measurement. Section 4 presents my findings on the effects of the treatments on class attendance and test scores. Section 5 concludes.

2 Theoretical Framework

This section presents a simple two-period model of household education production. In the first period, the household decides how much effort to exert in the absence of external incentives. In the second period, an experimenter augments the household's value of education with additional incentives to learn.

In this model, the education output produced in each period represents achievement in an individual educational task, such as grade promotion or exam performance. The parent and child each provide inputs, and the production process involves two-sided moral hazard similar to Eswaran and Kotwal (1985). The parent and child cannot contract on inputs, but they can contract on a division of the value that education produces. This value can take the form of either the parent's own value of education or an external incentive provided by the experimenter. Each agent provides input based on his or her share of the value generated.

In the first period the share of the value accruing to each agent is determined endogenously as part of the parent's maximization problem. The parent places a value on success in the task, but the child does not.[4] The child derives value only from the transfers he receives from his parent. The parent provides transfers to her child to induce effort, and the parent and child decide how much input to provide based on the surplus that each receives.

In the second period the experimenter offers an incentive for success either to the parent or to the child. If the incentive is given to the parent, the parent decides on a fraction of the total value (her own value plus the additional incentive) to give to the child, and the parties decide on second-period effort. If the incentive is given to the child, the parent keeps only her own value of success,

[4] I assume that the parent places a value on her child's education either directly through her utility function, or through future transfers from the child.

while the child receives the full amount of the additional incentive. This division determines the parties' second-period effort and the probability of success in the second period.

After describing the family's period-1 decisions in Section 2.1, I show how the family reacts to external incentives targeted to the parent or to the child in Section 2.2. In Section 2.3 I introduce commitment problems on the part of the parent and explore the implications of offering the parent a device that allows her to commit to giving her child full incentives.

2.1 Initial Period

The child and parent contribute c_t and p_t, in periods $t = \{1, 2\}$. The costs of input to the child and parent are given by the quadratic functions $\frac{ac_t^2}{2}$ and $\frac{bp_t^2}{2}$ in each period.[5]

The parent places a normalized value of 1 on success in each period. That is, she receives 1 if the task is successful, and 0 if it is not.[6] The probability of success in each period is given by

$$f(c_t, p_t) = \delta c_t + \beta p_t$$

The parameters $\delta \geq 0$ and $\beta \geq 0$ capture the the child's and parent's productivity in education production, respectively. Child and parental inputs reflect each agent's direct influence on education production. Child input represents the child's direct effort, while parental input represents the parent's own effort in assisting the child with studies or in directly managing the child's activities.

Before first-period input choices are made, the parent announces a fraction γ_1 of the value of success that she will give to the child if the task is successful. I assume that the parent must commit to this division of value until the first-period outcome has been realized.[7] I further assume that the parent cannot make negative transfers to the child. In equilibrium, this limited liability constraint will bind. If the child is unsuccessful, he will receive no transfer, and if he is successful, he will receive a fraction γ_1 of the value of success.

Based on the fraction γ_1, the child and parent maximize their respective shares of the value of

[5] I assume that the cost parameters a and b are sufficiently large that the probability of success is always less than 1.

[6] I assume that all parents have the same value of education in the absence of external incentives. If parents have different values, then a greater value of education implies that the restriction imposed by the child incentives treatment is less likely to bind. Therefore, the model predicts a smaller difference between parent and child incentives as the parent's value of education increases.

[7] This assumption will be relaxed in Section 2.3.

success net of costs over their first-period inputs. The first-order conditions of these maximization problems form two incentive-compatibility constraints that the parent faces in deciding on γ_1.

For simplicity I assume that both the child and parent are risk neutral and therefore make their decisions based on the expected value of success $f(c_1, p_1)$. The child's incentive-compatibility constraint is formed by the maximization of his share of the value of success net of costs over c_1, taking the parent's choice of γ_1 as given:

$$\max_{c_1} \gamma_1 (\delta c_1 + \beta p_1) - \frac{a c_1^2}{2}$$

This yields the first-order condition

$$c_1 = \frac{\delta \gamma_1}{a} \tag{1}$$

Similarly, the parent maximizes

$$\max_{p_1} (1 - \gamma_1) \{\delta c_1 + \beta p_1\} - \frac{b p_1^2}{2}$$

This yields the first-order condition

$$p_1 = \frac{\beta(1 - \gamma_1)}{b} \tag{2}$$

Subject to the incentive-compatibility constraints (1) and (2), the parent maximizes her expected share of the value of success over the fraction γ_1 she gives to the child:

$$\max_{\gamma_1} (1 - \gamma_1) \{\delta c_1 + \beta p_1\} - \frac{b p_1^2}{2}$$

Substituting the incentive-compatibility constraints, this optimization becomes

$$\max_{\gamma_1} (1 - \gamma_1) \left\{ \frac{\delta^2 \gamma_1}{a} + \frac{\beta^2 (1 - \gamma_1)}{b} \right\} - \frac{\beta^2 (1 - \gamma_1)^2}{2b}$$

The first-order condition for γ_1 is given by

$$\frac{(1 - \gamma_1)\delta^2}{a} - \frac{\delta^2 \gamma_1}{a} - \frac{\beta^2 (1 - \gamma_1)}{b} = 0 \tag{3}$$

Because $0 \leq \gamma_1 \leq 1$, The parent's choice of γ_1 will be

$$\gamma_1^* = \begin{cases} \frac{b\delta^2 - a\beta^2}{2b\delta^2 - a\beta^2} & \text{if } \beta < \delta\sqrt{\frac{b}{a}} \\ 0 & \text{if } \beta > \delta\sqrt{\frac{b}{a}} \end{cases} \tag{4}$$

Three important relationships arise from this optimization. First, the parent never gives the child full incentives (i.e., $\gamma_1^* = 1$), even when her input is completely unproductive. In the formulation above, $\gamma_1^* = \frac{1}{2}$ when $\beta = 0$. Second, over the range in which the solution is on the interior, γ_1^* is increasing in δ and decreasing in β. The parent will give higher incentives to the child as the child's input becomes more productive, and lower incentives to the child as the parent's own input becomes more productive. Third, the first-period probability of success (i.e., $\delta c_1^* + \beta p_1^*$) is increasing in both δ and β (see Appendix A for proofs).[8]

2.2 Experiment

Now suppose that after the first-period input decisions are made, an experimenter observes the expected output, $\delta c_1^* + \beta p_1^*$, and then offers an external incentive of value π to either the parent or child conditional on success in the second period. As noted above, the parent places a value of 1 on success in addition to this incentive.

If this external incentive accrues to the parent, the parent and child make their second-period production decisions based on the external incentive plus the parent's own value of success, $\pi + 1$. Based on this value and an updated share γ_2, the parent and child make their choices of c_2 and p_2. The child's choice of c_2 is determined by the maximization

$$\max_{c_2} (\pi + 1)\gamma_2 \left\{\delta c_2 + \beta p_2\right\} - \frac{ac_2^2}{2}$$

[8]The relationship between the probability of success and β does not depend on the assumption of quadratic costs. As β increases, there are three effects: 1) the probability of success increases directly, 2) the parent's incentive to contribute inputs increases, and 3) the parent shifts more resources towards herself through decreasing γ. The net effect is always positive, since by decreasing γ the parent is shifting resources to the more productive input. On the other hand, the relationship between the probability of success and δ is more dependent on the assumption of quadratic costs. While the probability of success still increases from the direct effect and from increased incentives for the child to contribute inputs, the parent may have an incentive to capture more of the surplus by *decreasing* the child's share γ. If this is the case, then it is possible that the parent's incentive to decrease γ overwhelms the direct effect and the effect on the child's incentives to contribute.

The first-order condition is given by

$$c_2 = \frac{\delta(\pi+1)\gamma_2}{a} \tag{5}$$

Similarly, given γ_2, the parent maximizes

$$\max_{p_2} (\pi+1)(1-\gamma_2)\left\{\delta c_2 + \beta p_2\right\} - \frac{bp_2^2}{2}$$

The first-order condition is given by

$$p_2 = \frac{\beta\left(\pi+1\right)\left(1-\gamma_2\right)}{b} \tag{6}$$

The parent's optimization over γ_2 is now

$$\max_{\gamma_2} \left(1-\gamma_2\right)\left(\pi+1\right)\left\{\delta c_2 + \beta p_2\right\} - \frac{bp_2^2}{2}$$

subject to the new incentive-compatibility constraints (5) and (6). The first-order condition is now

$$\frac{\left(\pi+1\right)^2\left(1-\gamma_2\right)\delta^2}{a} - (\pi+1)\delta^2\left[\frac{\left(\pi+1\right)\gamma_2}{a} + \beta^2\frac{\left(\pi+1\right)\left(1-\gamma_2\right)}{b}\right] = 0$$

Note that this condition is equivalent to equation (3), so that $\gamma_2^* = \gamma_1^*$.

Based on the share allocated to the child γ_2^* and the second-period inputs, the probability of success in the second period is given by

$$\delta^2\frac{(1+\pi)\gamma_2^*}{a} + \beta^2\frac{\left(1+\pi\right)\left(1-\gamma_2^*\right)}{b}$$

The probability of success, as before, is increasing in δ and in β.

Now suppose that the external incentive of π is given to the child and cannot be appropriated by the parent. This restriction binds if the parent would have given the child less than π if she had received the incentive herself. That is, the restriction binds if

$$(\pi+1)\gamma_2^* < \pi \tag{7}$$

I assume throughout that π is large enough such that the restriction (7) binds. If this is the case, then the parent will not offer the child any additional incentives; that is, if the task is successful she will not share the value that accrues to her. She will receive a value of 1 if the task is successful, and the child will receive a value of π. The probability of success in this case is given by

$$\delta^2 \frac{\pi}{a} + \frac{\beta^2}{b}$$

The difference in the second-period probability of success between incentives to the parent and incentives to the child is given by

$$\delta^2 \frac{(\pi+1)\gamma_2^*}{a} + \beta^2 \frac{(\pi+1)(1-\gamma_2^*)}{b} - \delta^2 \frac{\pi}{a} - \beta^2 \frac{1}{b} \tag{8}$$

I now explore how equation (8) is affected when either δ or β varies across households. Suppose first that δ varies and that β is fixed. At low values of δ, incentives provided to the parent will result in a higher probability of success than incentives provided to the child. At high values of δ, incentives provided to the parent will result in a lower probability of success than incentives provided to the child.[9]

The intuition behind this result is as follows. At very low values of δ, the child's input is unproductive, and the probability of success is maximized when the parent is allocated the incentives. As δ rises, the child's input becomes relatively more important in production, but because of moral hazard, the parent is inclined to distort the rewards towards herself. Therefore, at high values of δ, it becomes more effective to allocate π directly to the child.

Now suppose that β varies and that δ is fixed. An external incentive provided to the child will result in higher probability of success when β is low, and lower probability of success when β is high.[10] At low values of β, the probability of success is higher when the child receives high-powered incentives, but if given the option the parent distorts the rewards towards herself in order to gain

[9] Equation (8) will be negative when $\delta > \beta\sqrt{\frac{a}{b}}$ and positive when $\delta < \beta\sqrt{\frac{a}{b}}$. To see this, note that when $\delta = 0$, $\gamma^* = 0$ and equation (8) is positive and decreasing in δ. Equation (8) equals zero only when $\delta = \beta\sqrt{\frac{a}{b}}$, and at this point it has a negative partial derivative with respect to δ, implying that it is negative at higher values of δ.

[10] Equation (8) will be negative when $\beta < \delta\sqrt{\frac{b}{a}}$ and positive when $\beta > \delta\sqrt{\frac{b}{a}}$. To see this, note that when $\beta = 0$, $\gamma_2^* = \frac{1}{2}$ and equation (8) is negative. Equation (8) equals zero when $\beta = \delta\sqrt{\frac{b}{a}}$, and at this point it has a positive partial derivative with respect to β, implying that it is positive at higher values of β.

more surplus. For high values of β, it becomes more effective to provide incentives to the parent because she will allocate them to the more productive input (by keeping them for herself).

If child and parental productivity are unobservable, then the child's learning level at the beginning of the experiment, $\delta c_1^* + \beta p_1^*$, will be directly related to both productivity parameters and to the relative effectiveness of incentives given to the child. The share of the variation in initial test scores related to child productivity will be positively related to the relative effect of incentives given to the child. Conversely, the share of the variation in initial test scores related to parental productivity will be negatively related to the relative effect of incentives given to the child.

If households primarily differ by child productivity, then initial test scores will reflect differences in child productivity, and children with higher test scores will perform better in the experiment when they are given the incentives directly. If, on the other hand, households primarily differ by parental productivity, then children with lower test scores will perform better in the experiment when they are given the incentives directly. The following table illustrates the relationship between relative treatment effects and the children's initial test scores when the variation in these scores is driven by differences in child or parental productivity.

		Variation in Productivity	
		Child (δ)	*Parent (β)*
Initial Test Score	*Low*	Parent incentives > Child Incentives	Child Incentives > Parent Incentives
	High	Child Incentives > Parent Incentives	Parent incentives > Child Incentives

2.3 Parental Commitment

Now suppose that the parent cannot commit to rewarding her child after the outcome has been realized. In this case, the parent will not reward the child at all, effectively setting $\gamma = 0$. If the child is aware that the parent cannot commit, he will not provide any input. Thus, the probability of success in the first period will be $\beta p_1 = \frac{\beta^2}{b}$.

If an external incentive of π is provided to the parent and she cannot commit, the probability of success in the second period equals $\beta^2 \frac{\pi+1}{b}$. The parent will desire a commitment device that forces her to give the external incentive of π to the child if

$$\delta c_{2,commit} + \beta p_{2,commit} > (\pi + 1)\beta p_{2,nocommit} \tag{9}$$

Using the cost functions from above, this condition is equivalent to

$$\beta^2 < \delta^2 \frac{b}{a} \frac{1}{2+\pi}$$

Note that this condition will hold when β is low, or when δ is high. When the parent's input is not productive (low β), the parent will want to commit because she does not lose much in terms of her own contribution to success by committing. Similarly, when the child's input is very productive (high δ), the parent will want to commit because she stands to gain more (in terms of increased likelihood of success) by doing so.

In addition, the parent will only be willing to commit when doing so increases the probability of success. To see this, note that equation (9) implies that the second-period probability of success is greater when the parent decides to commit:

$$\delta c_{2,commit} + \beta p_{2,commit} > \beta p_{2,nocommit}$$

Thus, allowing the parent to commit will increase the probability of success relative to an unrestricted parent incentive if commitment is indeed a problem.

3 Experimental Design

The intervention was conducted from July through September of 2007 in Gurgaon, a suburb of Delhi. Eight government-run primary schools were selected based on proximity to the city center and availability of public transportation nearby. In seven schools all 1st, 2nd and 3rd grade students participated in the program. In one school administrative difficulties in obtaining 1st-grade students' addresses prevented us from including 1st-grade children.[11]

The main intervention consisted of a pretest, announcement of the child's incentive scheme, and a post-test approximately two months later. Children were initially tested in schools to determine baseline learning levels. Each child was given a goal competency based on his pretest score. Children were then re-tested using a similar testing instrument after two months. If the child

[11] The sample also includes 36 students from two additional schools near the city center. All of the results are robust to the exclusion of these students.

achieved the goal, he or his parent would receive a prize. The treatments were randomly assigned on an individual level (within schools), after stratifying by pretest score.

The prize value was set at 100 rupees (about $2.50 at the prevailing exchange rate) for all treatments. One hundred rupees is the approximate daily wage for an unskilled laborer in these areas. This amount is substantially less than the amount of $20 awarded in the program evaluated by Kremer, Miguel and Thornton (2004) in Kenya. However, there are several reasons to believe that the relatively smaller amount in the India program was sufficient to induce noticeable changes in behavior. First, this study's time frame of 2 months was considerably shorter than the Kenya program, which ran for 8-11 months from announcement to post-test. Second, the Kenya program was structured as a tournament that only awarded 15% of the girls with prizes. The India program's goals were set so that everyone in the program had approximately a 50% likelihood of winning the prize, calibrated by a pilot experiment. Third, the individual announcement of the program at the children's homes and the two individual reminders should have made the program more salient to the parents and children who participated. Finally, in Section 4.1 I offer more direct evidence that the prize value was sufficient to raise test scores. Using households that were not reached for the program announcement as controls, I find that the program raised test scores by about 0.51 standard deviations across all incentive treatments.

The experiment exogenously changed the recipient of the reward along two dimensions: the direct recipient of the reward (either parent or child) and type of reward (either money or toys). The first treatment ("parent money") offered a reward of money to the parent if the child achieved the goal. The parent was visited at home within three days of the post-test and was given 100 rupees in cash. This represents the case where the parent has complete control over the reward. The second treatment ("child money") offered a reward of money to the child if the child achieved the goal. The child was given 100 rupees in cash at school within three days of the post-test. If the parent and child consider money given to the child as earmarked for child consumption and there is no compensating behavior by the parent, this treatment would represent more control over the reward by the child. However, if income from the child and parent are pooled in the household, this treatment would be equivalent to the parent money treatment.

Several responses from the baseline survey suggest that in my sample money given to the child was pooled within the household. While 51% of parents reported giving their child spending money over the past day, the majority of the time (79%) the money was given for school supplies or food items. This suggests that money was rarely given to the child to be spent at his own discretion. The baseline survey also asked the child what he would do if given 100 rupees. Over 80% reported that they would give it directly to their parent.

The third treatment offered a reward of a toy valued at 100 rupees to the child if the child achieved the goal. There were two sub-treatments. In the first sub-treatment ("toys"), the child was given a choice between a menu of 5 different toys, each with a retail value of 100 rupees. In the second sub-treatment ("voucher"), the child was given a voucher worth 100 rupees that was redeemable at a local toy store. The voucher sub-treatment was included in case the toys chosen for the toys treatment were unpopular. This would allow some children to receive an item they valued more. In practice, however, the limited number of toys selected for the toys treatment proved to be very popular, and the shopkeepers reported that those redeeming the voucher often chose toys that were available in the toys treatment. Therefore, these sub-treatments are combined for the purposes of the analysis.

Rewarding the child with a toy gave him control over the reward in two ways. As in the case of the child money treatment, it was given directly to the child. In contrast with the child money treatment, it also gave the child an item that could not easily be used by the parent. While the parent still retained the right to take away the toy from the child, it would have been difficult to sell the toy and convert its value to other household consumption. The value of the toy was also high enough that the parents were unlikely to be able to adjust the child's consumption of these goods, at least in the short term. While 4% of parents reported having given their child a toy during the week before the baseline survey, anecdotal evidence suggests that the value of these toys was substantially less than the toys offered as part of the program.

In order to test the commitment hypothesis, I included two additional treatments that offered the parents a choice between money for themselves and a toy for their child. In *ex ante choice* treatment, the parent made her choice at the time of program announcement. In the *ex post choice* treatment, the parent made her choice after the child had reached the goal.

As shown in the previous section, a parent who is unable to commit to a division of surplus

γ will desire to yield control to the child *ex ante* if the inequality (9) holds. These parents are expected to choose the toy reward in the *ex ante* treatment. Lack of commitment also implies that outcomes in the *ex ante* choice treatment will be higher than outcomes in the money treatments.

The *ex post* choice treatment was included to confirm that the salience or the convenience of the choice itself does not cause a positive impact of the *ex ante* choice. If the choice itself drives the results, one should expect a positive impact of either choice treatment on outcomes. On the other hand, if the results are driven by an actual desire to commit, only the *ex ante* choice treatment will positively affect outcomes. In addition, the *ex ante* and *ex post* choice treatments can serve to check for consistency in choices. If parents desire to commit *ex ante* because they know they will not reward the child *ex post*, they will choose money *ex post*.

Regardless of treatment category, all children were invited to attend free after-school classes run as part of the program. The classes were led by teachers trained to assist the children in achieving their literacy goals. The profile and training of the teachers followed Pratham NGO's "balsakhi" model (see Banerjee et. al., 2007). In each school, enough teachers were provided so that there was at least one teacher for each 20-30 students who attended the classes.[12] Classes ran for three hours every afternoon that school was in session.[13] Children were free to attend on a drop-in basis, and teachers were given flexibility to customize lessons based on the reading levels of the children who attended. Tutorials held outside of school hours are common in India, and thus the extra classes provided a learning environment familiar to the households in the study.

There were two primary reasons for including the classes. First, government schools in India are often poor platforms for learning, and the classes provided a greater chance for the children to reach the goals set by the program. Second, the classes present a unique opportunity to measure effort that is not present in most studies of education interventions. Existing studies traditionally rely on attendance in school, taken from either 1) the school's administrative records, or 2) random, unannounced checks by outside surveyors. Administrative records are notoriously inaccurate in India, as schools may have incentives to inflate attendance (Linden and Shastry, 2008). Random checks are usually unable to measure attendance on a daily basis, since they disrupt the classes and

[12] In one school, the principal did not allow our teachers to access the school premises to conduct the class, and no suitable alternative location was found.

[13] Schools in Gurgaon run from 8am to 12:30pm.

are difficult to take accurately. Extra-class teachers in this study were familiar with the students so that the attendance records could be easily matched with the survey data.

The nature of the decision to attend these classes may reflect both parent and child effort. The parent can contribute by reminding and encouraging the child to attend, and the child can contribute by taking the initiative to attend on his own. If this effort is a reflection of the overall production function, then the theoretical predictions for class attendance mirror those for achievement. In addition, because effort does not include the stochastic component inherent in achievement of the goal, the results are expected to be stronger.

3.1 Pretest

Children were initially tested for reading ability during school time. The test used an instrument developed by Pratham, a large India-wide NGO specializing in child literacy. The test evaluated each child on a four-point scale: 0) the child could not recognize letters, 1) the child could recognize letters, 2) the child could read simple words, 3) the child could read a simple paragraph, and 4) the child could read and understand a several-paragraph story.

Based on each child's ability at the time of the pretest, the child was given a goal competency to be reached when he was re-tested after two months. If the child achieved the goal, he or his parent would receive the prize as per his treatment category. Children reading at levels 0, 1, and 3 were each given a goal one level above their current competency, while children at level 2 at the pretest were given a goal of 4. Goals were selected such that approximately half of the children would reach the goal, based on a pilot study.

Each child was then randomly assigned to one of the six treatment groups. The randomization was conducted at the individual level and was stratified by pretest score within each school, grade and classroom. Children at the highest reading level (4) at the pretest were excluded from the study and were instead given an unconditional prize at the end of program.

3.2 Program Announcement / Baseline Survey

Approximately one week after the pretest, a baseline survey was conducted at the child's home, and the incentive program was announced to the mother and child. The survey and the program announcement were conducted with the child's mother rather than his father because pilot surveys indicated that the mother was usually more involved in the child's education than the child's father. The survey collected demographic information as well as information on the transfers that the parent had made to her child over the past week.

After finishing the survey, the surveyor read a script announcing the incentives program to the mother and child. The script was individualized based on each child's treatment group and baseline test score. The mother and child were informed that the child would be tested again in school after two months, and if the goal competency was reached, the child or parent would receive the specified prize. In addition to the announcement of the incentive scheme, the mother and child were informed that special classes would be conducted from 2 to 5pm after school in order to assist the child in reaching the goal.

Table 1 shows the sample composition by pretest score, grade and treatment category. Out of 1466 children who took the pretest, 331 were excluded from the study because they achieved the highest possible test score, and 49 others were excluded because they lived too far from the schools, making surveying impractical. 1086 children were thus available for the randomization. 85% of children out of the randomized group of 1086 were reached for the baseline survey and program announcement. These 925 children form the sample who received the treatment. The attrition between the randomization and program announcement was primarily due to the difficulty in locating the children's homes and in availability of parents during the day.[14] Of the 925 children offered program, 900 (97%) took the post-test after two months. Most of the 25 students who were not available for the post-test had moved away since the program announcement. Attrition between the program announcement and post-test is not significantly related to either treatment category or test score. Out of the 11 baseline variables in Table 2, attrition is significantly positively related to total household members (p-value = 0.082), but is not significantly related to any of the other variables (results not shown). Given the large number of variables examined and the general lack of significance, it is reasonable to conclude that attrition does not bias the achievement results.

[14] Schools in Gurgaon do not keep detailed information on the addresses of their students. It was therefore necessary to have every child in the study show the surveyor his home at the time of the pretest, so that the surveyor could note the child's address information. In some cases, the children were not available to show the surveyors their homes, they could not find their homes (as in the cases where parents or older siblings brought them to school), or their homes were prohibitively far from the school that conducting three surveys there was impractical.

The attendance results include all children offered the program regardless of whether they took the post-test.

Table 2 confirms the effectiveness of the randomization by examining correlations between treatment status and baseline household characteristics and behaviors.[15] In Panel A, each column regresses the relevant characteristic or behavior on dummies indicating each treatment category, omitting the parent money category. Panel B repeats this exercise, comparing only the two toys categories with the two money categories.

Among the six treatment categories, several observables are significantly related to treatment status, but this is not surprising given the large number of regression coefficients in the table. In Panel B, the broader treatment categories are marginally related to 1 variable out of the 11 tested. While Table 2 shows that treatment status is not systematically related to observables, I include all of these variables as controls in my outcome specifications to confirm that the results are robust to these controls. As I show in the next section, inclusion of these controls does not affect the results in any meaningful way.

Table 3 contains summary statistics of household demographics and incentives that the parent gave to the child at the time of the baseline survey. On average, parents gave their children 13 rupees per week for spending money; in the majority of cases, this was given for the child to purchase food. As noted earlier, 4% of parents reported having given their children toys over the past week, although it is likely that the value of these toys were small relative to those offered as part of this program. A number of parents (9%) also reported having given their child an item other than toys over the past week. These items usually consisted of school supplies or clothes.

3.3 Follow-up Survey/Reminder

Surveyors returned to the households approximately one month after the announcement of the program to conduct another survey that collected information on transfers between parents and children. At the end of the survey, a short script was read reminding the parent and child of the program and specifying the date of the post-test. Households were given cards that contained

[15] The mother and father education variables were mistakenly excluded from the baseline survey and had to be measured at the second followup. Since these are objective measures, however, it is unlikely that survey responses were biased by the treatments.

information on the child's goal, prize and test date.

3.4 Post-test

Approximately two months after the program announcement, a post-test, similar in form but not exact content to the pretest, was conducted in the schools. Prize distribution was conducted the day after the post-test either at school or at the child's home, according to the child's treatment category. Prizes for the child money, toy and voucher treatments were distributed in school, in addition to toys chosen in the *ex ante* treatment. Parents in the parent money treatment and those who chose money in the *ex ante* treatment were given the money at home. Parents in the *ex post* choice treatment were also visited at home the day after the post-test and were given their prizes upon making the choice.

3.5 Second Follow-up Survey

Approximately one week after the post-test, a second follow-up survey was conducted to again measure transfers given by the parent to her child. The purpose of this survey was to examine how parents reacted after the post-test had been conducted and rewards had been distributed.

4 Results

4.1 Overall Effect of the Incentives Program

The effects of incentives programs are important from a policy perspective, and it is useful to show that the incentives used in this study did influence test scores overall. The study design did not include a pure control group, and therefore I cannot offer experimental evidence on the overall effects of giving children incentives compared to not giving them incentives.[16] I can, however, offer non-experimental evidence of these overall program effects.

As shown in Table 1, a number of children and their mothers in the randomized sample were not reached at the time of the baseline survey. Some children were not in school when addresses were initially collected, but they were included in the randomization in case they could be found at

[16] In order to gain support of the local school committee, the study was designed to treat all eligible first, second and third graders in the program schools.

a later date. For others, the children and their mothers were not available during the surveyors' 2-3 daytime visits when the baseline was conducted. There were 161 students included in the randomization but who were not reached for the program announcement. Out of these 161 students, 152 were in school when the post-test was administered. This group (the "no program" group) will serve as a control group for the analysis of this section. The results in this section include both children whose addresses were not collected and those who were not available at the time of survey, but restricting the sample to either of these groups leaves the results largely unchanged.

One caveat is in order with respect to interpretation of the results presented in this section. While the after-school classes were open to any child who wished to attend, children in the incentives program were notified individually when the program was announced. In practice, children in both the program and no-program groups were often reminded of the classes during school time by their teachers, but this was not controlled as part of the experiment. Therefore, the program effects estimated in this section are the combined effects of receiving an incentive treatment in addition to individual notification of the classes. Children not in an incentive treatment did attend the classes, however, at rates of approximately half as much as those in an incentive treatment. It is not possible to determine if these differences are driven by the individual notifications of the classes, or by the additional motivation provided by the incentives themselves.

Figure 1 presents the distributions of raw pretest and post-test scores of the program group and the no-program group. As shown in the top panel, the pretest scores are remarkably similar between the two groups. A Pearson χ^2 test fails to reject the equality of the two distributions (p-value = 0.89). The bottom panel presents the post-test scores of the two groups. The program group now has a much lower proportion of test scores of zero and higher proportions of test scores of 2 and 4. There is a lower proportion of scores of 3, but this result is not surprising given that this score category was not one of the goals given to students in the program group. A Pearson χ^2 test now strongly rejects equality of the two distributions (p-value < 0.01).

Table 4 presents the results of regressions of a dummy variable that indicates if the mother and child were reached at the baseline on relative pretest score and other observable characteristics. Column 1 indicates that being reached at the baseline is not significantly related to pretest score, grade or the gender of the child. Columns 2 and 3 add controls for classroom and surveyor dummies. In the specification with both sets of dummies, a one standard-deviation increase in pretest score is related to a 2.9% *lower* likelihood of being reached at the baseline. Part of this relationship may be explained by the availability of the children: if higher-performing children attend tutorials outside of school time, they will be more difficult to reach.[17]

Table 5 presents the difference-in-difference estimates of participation in the program on normalized test scores.[18] In Column 1, I present the regression without controlling for observable characteristics. The result is large and highly significant: program participation is associated with a 0.51 standard deviation increase in test scores relative to children whose parents were not reached for the baseline survey. The precision of the estimate increases when additional controls are added in Columns 2 and 3, but the estimate remains unchanged because none of these controls are time-varying.

On average, the no-program group increased by 0.15 standard deviations during the two months between the pretest and post-test. To check whether this is a reasonable increase in test scores in the absence of the program, I compare this increase to the cross-sectional differences in test scores between grades 1, 2 and 3. On average, a first grader is 0.7 standard deviations below a second grader at the pretest, and a second grader is 0.7 standard deviations below a third grader. Gurgaon schools are in session 11 months out of the year, and if a student improves by an equal amount each month, we would expect a student to improve $2/11 * 0.7 = 0.13$ standard deviations on average. This is very close to the 0.15 standard deviation increase observed among the no-progarm group.

4.2 Main Treatments

Although the model does not provide a prediction on the differences between incentives to parents and incentives to children overall, it is informative to start with these differences before turning to the interactions.[19] The first three columns of Table 6 show these differences using attendance in the after-school classes as the outcome. The dependent variable is a dummy indicating any

[17] In the specification with classroom dummies only, the female dummy is associated with a 9% lower likelihood of being reached. Because classes are often split by gender, the specification with classroom dummies removes a significant portion of the variation in the female dummy, making the coefficient difficult to interpret.

[18] Test scores were normalized based on the mean and standard deviation of all pretest scores in the child's grade.

[19] The results of the choice treatments are reported and discussed in Section 4.3.

attendance in the after-school classes over the two-month period in which they were held.[20] In all three regressions, the differences between treatment categories are small and insignificant. Columns 4 through 6 present the results of similar regressions using achievement of the goal competency as the outcome variable. As with the attendance results, the differences between treatment categories are small and insignificant.

As noted earlier, survey responses indicated that most parents have discretion over money received by their children. In order to maximize power, I pool the two money treatments and compare them with the toys treatment (including the voucher sub-treatment) in analyzing interactions. The money treatments therefore represent parental incentives, and the toys treatment represents child incentives.

The model predicts that child incentives will be more effective in encouraging attendance and achievement when the productivity of the parent's input is low, and parental incentives will be more effective when the productivity of the parent's input is high. It also predicts that child incentives will be more effective when the productivity of the child's input is high, and parental incentives will be more effective when the productivity of the child's input is low. Child and parental productivity in education production are not directly observable, but the child's initial test score may reflect parent or child productivity, depending on the primary source of variation in test scores. If the variation in test scores is driven by differences in parental productivity, then the model predicts a negative relationship between the child's test score and the relative effectiveness of child incentives. If test scores reflect differences in child productivity, then the model predicts a positive relationship between the child's test score and the relative effectiveness of child incentives.

While parental and child productivity are not directly observable, it is possible to use parent characteristics, as measured in the baseline survey, as proxies for parental productivity. As described in the theory section, parental productivity represents the extent to which the parent can either provide direct inputs into the child's education or manage a child's schooling behavior. For example, parents can spend time helping their children with schoolwork, or they can pay for the child to receive outside tutoring. In addition, parents with more time at home will be more able to manage their children's study and schooling behavior directly.

[20] All regressions with control variables include dummies for missing values. Restricting the sample to observations with no missing values of these variables does not stubstantially affect the results.

I include eight variables from the baseline survey as proxies for parental productivity. These variables fall into four broad categories. First, I include two variables that reflect household composition. More children below age 15 in the household should take the parent's time away from the program child and therefore are expected to negatively affect the parent's ability to contribute to the child's education. On the other hand, the number of household members at or above age 15 are expected to positively affect productivity, since these members represent resources the child can use for help with his studies. Second, I include three variables that reflect the work and education status of the child's parents. The mother's employment status could affect her ability to contribute to her child's education because employed mothers will have less time to devote to their children.[21] Mother and father's education are also included because more educated parents should increase the parent's ability to help the child with studies. Third, I include durables ownership, a measure of household wealth. Household wealth is expected to be positively related to parental productivity because parents in more wealthy households spend less time meeting basic needs and can therefore devote more time to their children. In addition, more wealthy households can contribute resources such as school supplies to facilitate their children's education. Finally, I include two measures of productive behavior: an indicator for whether anyone in the household has helped the child with his studies over the past day, and the total amount of money that the parent spent on tutoring for the child over the past month.

To capture the extent to which these variables influence child test scores, I first regress the child's relative pretest score on these variables. I construct the student's relative pretest score as the difference between the student's own pretest score and the grade-specific mean pretest score, normalized by the grade-specific standard deviation in test score.[22] The results of the first-stage regression of relative pretest score on the parental productivity variables are reported in Appendix Table 1.[23] In the main specification of Column 1, all of the estimates except for father's education and durables ownership have the expected sign, and of the six with the expected sign four are

[21] Father's employment status is not included because only 4% of fathers were reported to be unemployed. It is also not clear *ex ante* whether father's unemployment represents additional parental resources, or if the father is unemployed because he is sick, injured, or otherwise incapacitated.

[22] In order for the grade-specific means to be representative of all children in the schools, I include the highest pretest scores in constructing this variable.

[23] Because households with children whose pretest scores were in the highest category were not surveyed, these scores are not included in this regression.

significant at the 10% level.[24]

I then use the predicted values of the child's relative pretest score in the following regression:

$$y_i = \alpha_0 + \alpha_1 toys + \alpha_2 zpretest_i + \alpha_3 toys_i * zpretest_i + \kappa' X_i + \varepsilon_i \qquad (10)$$

where y_i represents the outcome for child i (either attendance or achievement of the literacy goal), $toys_i$ is an indicator for whether child i was in the toys treatment, $zpretest_i$ is the child's relative pretest score and X_i is a vector of controls.

The first three columns of Table 7 present OLS estimates of equation (10) using the actual values of *zpretest*. These regressions use attendance in the after-school classes as the outcome of interest. The estimated coefficients on the interaction terms are negative and significant in all three specifications, implying that the toys treatment resulted in higher attendance relative to the money treatments among children with low pretest scores and lower attendance among children with high pretest scores. The magnitude of the coefficients imply that an increase in one standard deviation of the child's pretest score results in a 12-13% decrease in the likelihood of attending class if the child is in the toys treatment.

Columns 4 and 5 of Table 7 present estimates of equation (10) using predicted values of *zpretest*, based on the regression in Column 1 of Appendix Table 1.[25] The estimated coefficients on interaction terms are negative and significant at the 5% level in both specifications. The magnitudes of these coefficients are more than three times the estimates using the actual values of *zpretest*, suggesting that the relative effectiveness of the toys treatment is more strongly related to the portion of pretest scores driven by parental characteristics than to pretest scores overall. The difference in magnitudes is consistent with the theory, since the theory predicts that relative effectiveness of the toys treatment will be directly related to the *share* of test scores predicted by the parent's contribution.

Table 8 repeats the analysis of Table 7 using achievement of the literacy goal as the outcome. The first three columns present estimates of equation (10), using the actual values of *zpretest*.

[24] It is somewhat surprising that durables ownership is significantly related to *lower* initial achievement. Note, however, that this relationship is driven in part by tv ownership, which could lower the child's motivation to study.

[25] Because $zpretest_i$ is generated from a first-stage regression, standard errors for these regressions are bootstrapped based on 500 replications.

Again, there is a negative relationship between the effect of the toys treatment and the child's pretest score, although the estimated coefficient is significant at the 10% level in only one of the three specifications. In Columns 4 and 5, I present the estimates using the predicted values of *zpretest*, similar to the corresponding columns of Table 7. As with the attendance results, the magnitudes of the estimated coefficients on the interaction terms increase considerably, although these estimates are not significant.

If proxies for child productivity were available, a similar analysis should yield a positive interaction term between the variation in pretest scores attributable to child productivity and the relative effectiveness of the toys treatments. Since I am not able to observe proxies for child productivity, I supplement the results above by using the variation in test scores driven by school and classroom characteristics. School and classroom characteristics should not be systematically related to either parent or child productivity, which implies that the coefficients on the interaction terms should be zero. I find that repeating the exercise of Tables 7 and 8 using the variation in test scores predicted by classroom characteristics yields small and insignificant coefficients on the interaction terms (results not shown).

Figures 2 and 3 present less parametric estimations of the treatment effects by initial learning level. Figure 2 graphs the estimated differences in attendance between the toys and money treatments by the student's raw baseline test score. The specification controls for classroom dummies, assuming these effects are constant across different baseline test scores. The treatment effects are clearly decreasing across all 4 categories of test scores. In the lowest score category, children in the toys treatment are 16% more likely to attend the classes, a difference significant at the 5% level. In the highest two score categories, children in the money treatments are 18% and 24% more likely to attend the classes for baseline scores of 2 and 3, respectively. Both of these estimates are significant at the 10% level. The F-test for the joint significance of all 4 effects definitively rejects the null hypothesis that all effects are zero (p-value = 0.013).

Figure 3 repeats this exercise using achievement of the goal competency as the outcome of interest. As with the effects on attendance, there is a monotonically decreasing relationship between the relative difference and pretest score. Children with the lowest pretest scores are 7% more likely to achieve the goal when in the toys treatment relative to the money treatments. In the highest test score category children are 18% less likely to achieve the goal when in the toys treatment relative to the money treatments. In this case, however, none of the individual effects are significant, and the F-test does not reject the null that all affects are zero (p-value = 0.39).

4.3 Choice Treatments

As described in the model, the incentives that the parent gives to the child may be dampened by both moral hazard and by an inability to commit to rewarding the child for positive outcomes. The *ex ante* and *ex post* choice treatments were included to determine whether commitment is indeed a problem.

Before turning to outcomes, I analyze the choices themselves. As implied by the model, if parents cannot commit and are aware of this problem, they would be willing to reward their children with a toy *ex ante* but would decide to keep the money for themselves *ex post*. Therefore, if the choice of a toy is primarily driven by a desire for commitment, one would expect more parents to choose the toy in the *ex ante* treatment than the *ex post* treatment (this behavior would constitute a preference reversal).

Panel A of Table 9 displays the percentage of parents in each treatment who chose toys. Substantially *more* parents in the *ex post* treatment chose to reward their children with toys: 33% of parents in the *ex ante* treatment and 51% of parents in the *ex post* treatment chose toys. Note, however, that the *ex post* choice reflects the selected group of parents whose children achieved the goal. To account for this selection effect, I compare the parents in the *ex ante* treatment whose children achieved the goal with those in the *ex post* treatment. The percentage of parents of achievers in the *ex ante* choosing toy remains virtually unchanged. Still, I cannot rule out the possibility that the differences are based on selection, so these results are merely suggestive. However, the fact that such a high number of parents chose to reward their children with a toy *ex post* provides some evidence against commitment problems. Instead, it is possible that parents initially chose money because they were uncertain of their needs for cash at the end of the program. Once the uncertainty was resolved, they were willing to choose the toy.

Recall that the model predicts that if commitment was a problem parents were aware of, lower productivity parents would be more likely to choose the toy *ex ante*. Panel B of Table 9 shows the results of a regression of the choice of toy on the parental productivity index, where this index is constructed as the predicted values of *zpretest*, as described in the previous section. In all three specifications, the coefficient on the productivity index is small, positive and insignificant. Thus, there is no evidence that low-productivity parents are more likely to choose the toy.

Turning to the effects of the choice treatments on outcomes, Columns 1-3 of Table 10 show the class attendance levels of children in each choice treatment. In the model, allowing parents to commit *ex ante* will improve outcomes if they do have commitment problems. Column 1 shows that there is no significant difference between attendance of children in the *ex ante* treatment compared with the money and toys treatments. Because attendance in the *ex post* treatment is also similar to that of the money treatment, the introduction of a choice itself does not appear to influence attendance. Column 2 includes interactions of each treatment group and the child's relative pretest score. If, as implied by the previous results, differences in test scores were driven by differences in parental productivity, commitment would have been the most helpful to parents of children with the lowest initial test scores. While the coefficient on the interaction term for the *ex ante* choice treatment is indeed negative, it is small in magnitude and not at all significant. Column 3 presents the interactions of the productivity index and the toys treatments. In this case, the coefficient on the on the *ex ante* choice interaction is large and *positive*, but the estimate is not at all significant.

Columns 4-6 of Table 10 repeat the analysis using the binary indicator for achievement of the literacy goal as the outcome of interest. The results are broadly consistent with those for attendance. The small and insignificant coefficient on the *ex ante* treatment dummy in Column 4 suggests that the option to commit to a toy at the beginning of the program does not increase achievement. Column 5 includes an interaction of the treatment groups and the child's relative pretest score. Again, there are no significant differences between the choice treatments and the money treatments at different pretest scores. Column 6 interacts the *ex ante* treatment with the productivity index, and the estimate is small, positive and insignificant. On balance, the results of the choice treatments present little evidence supporting the hypothesis that parents cannot commit to rewarding their children.

4.4 Survey Results

Throughout this paper I have assumed that parents must provide transfers to their children in order to induce them to learn. By using self-reported transfers, I can test this assumption. If parents do provide these transfers, then parents in the money treatments should provide more transfers to their children than those in the toys treatments. Further, if effort is unobserved, the transfers should occur primarily after the goal has been reached.

Table 11 shows presents differences-in-differences estimates of the effects of the toys treatment on transfers from parents to children, controlling for the differences in transfers measured in the baseline survey. The behaviors examined are the amount of money given to the child over the past week, a dummy for whether the parent gave a toy to the child over the past week and a dummy for whether the parent gave the child any other item over the past week. Panel A uses data from the first follow-up survey, taken just before the post-test. Column 1 uses an aggregate measure of the three types of transfers by averaging the z-scores of the measures, using the mean and standard deviation of each variable at the baseline survey as a base. The coefficient on the interaction term is small and insignificant, indicating that parents overall did not give more transfers in the money treatments prior to the post-test. Columns 2-4 report the difference-in-difference estimates for each type of transfer. The estimates are small and insignificant, except for the difference in non-toy items, which is significant at the 10% level.

Panel B of Table 11 repeats the difference-in-differences exercise using the differences between the second follow-up survey and the baseline survey. Using the aggregate measure of transfers in Column 1, the coefficient on the interaction term is large and significant, indicating that parents in the money treatment groups did provide more transfers after the outcome of the test had been realized. The treatment effect on giving other items is large in magnitude and highly significant, while the effects on money and toys given are small and insignificant. The estimated coefficient in the regression using other items as the dependent variable implies that parents in the toys groups gave their children other items 28% less often than those in the money groups. These results suggest that parents who received money as a prize often used this money to purchase more practical items such as clothes and books for their children. While the money was often spent on items used by the children, it is likely that it was not spent in the manner most preferred by the children themselves.

5 Conclusion

In this paper, I present the results of a field experiment in Gurgaon, India, designed to test the effects of changing the recipient of incentives to learn from parent to child. The experiment offered incentives for families of first, second and third graders in government primary schools to increase their children's reading ability. Consistent with a two-sided moral hazard model of education production where the variation in test scores is driven by differences in parental productivity, I find that for children with less productive parents and lower initial test scores, offering toys to the child as a reward results in better outcomes than offering money to the parent or child. On the other hand, for children with more productive parents and higher initial test scores, offering money to the parent or child results in better outcomes than offering toys to the child. These results hold for both attendance in after-school classes run as part of the program and for achievement of the target reading competency. By including a treatment that allowed the parent a choice between money and toys, I am able to provide evidence that these differences are not driven by the parent's inability to commit to providing incentives to her child.

Thus far the debate on cash transfer programs in developing countries has focused on which adult within the household should receive the transfer rather than including the possibility of a child receiving the transfer directly. In the case of education, outcomes depend on child effort, and schemes that reward parents must depend on the parents to both effectively exert their own effort *and* motivate their children. My results suggest that when the parent's productivity in education production is low, providing incentives to learn directly to the child may be more effective than providing incentives to the parent. While this paper has focused on children early in the education process, future research should examine decision making between older children and their parents to understand how the decision process changes as children grow older.

This paper also has implications for understanding how households respond to other external factors such as changes in returns to education. If parents' perceived returns to education change, households with more productive parents may respond more strongly because these households suffer from fewer agency problems between the parents and children. Since households with more productive parents tend to have higher-performing children to begin with, increases in returns may exacerbate inequality in education outcomes. This paper provides a starting point for future research examining these effects.

References

[1] Joshua Angrist, Daniel Lang, and Philip Oreopoulos. Lead them to water and pay them to drink: An experiment with services and incentives for college achievement. *NBER Working Paper 12790*, 2006.

[2] Joshua Angrist and Victor Lavy. The effect of high school matriculation awards: Evidence from randomized trials. *NBER Working Paper 9389*, 2002.

[3] Abhijit Banerjee, Shawn Cole, Esther Duflo, and Leigh Linden. Remedying education: Evidence from two randomized experiments in India. *Quarterly Journal of Economics*, 122(3):1235–1264, 2007.

[4] Felipe Barrera-Osorio, Marianne Bertrand, Leigh Linden, and Francisco Perez-Calle. Conditional cash transfers in education: Design features, peer and sibling effects: Evidence from a randomized experiment in Colombia. *NBER Working Paper 13890*, 2008.

[5] Gary Becker. A theory of social interactions. *Journal of Political Economy*, 82(6):1063–1093, 1974.

[6] Theodore Bergstrom. A fresh look at the rotten kid theorem-and other household mysteries. *Journal of Political Economy*, 97(5):1138–1159, 1989.

[7] Esther Duflo. Grandmothers and granddaughters: Old-age pensions and intrahousehold allocation in South Africa. *World Bank Economic Review*, 17(1):1–25, 2003.

[8] Mukesh Eswaran and Ashok Kotwal. A theory of contractual structure in agriculture. *American Economic Review*, 75(3):352–367, 1985.

[9] Roberta Gatti. Family altruism and incentives. *Scandinavian Journal of Economics*, 107(1):67–81, 2005.

[10] Paul Glewwe and Pedro Olinto. Evaluating the impact of conditional cash transfers on schooling: An experimental analysis of Honduras' PRAF program. *Mimeo, University of Minnesota*, 2004.

[11] C. Kirabo Jackson. A little now for a lot later: a look at the Texas Advanced Placement Incentive Program. *Mimeo, Cornell University*, 2008.

[12] Michael Kremer, Edward Miguel, and Rebecca Thornton. Incentives to learn. *NBER Working Paper 10971*, 2004.

[13] Leigh Linden and Kartini Shastry. Identifying agent discretion: Exaggerating student attendance in response to a conditional school nutrition program. *Mimeo, Columbia University*, 2008.

[14] Shelly Lundberg, Robert Pollak, and Terence Wales. Do husbands and wives pool their resources? Evidence from the United Kingdom child benefit. *Journal of Human Resources*, 32(3):463–480, 1997.

[15] Pratham Education Initiative. *Annual Status of Education Report*. Pratham Resource Center: Mumbai, 2008.

[16] Norbert Schady and Maria Araujo. Cash transfers, conditions, school enrollment, and child work: Evidence from a randomized experiment in Ecuador. *World Bank Policy Research Working Paper 3930*, 2006.

[17] T. Paul Schultz. School subsidies for the poor: Evaluating the Mexican Progresa poverty program. *Journal of Development Economics*, 74(1):199–250, 2004.

[18] UNESCO. *2008 EFA Global Monitoring Report*. Oxford, UK: Oxford University Press, 2007.

[19] Bruce Weinberg. An incentive model of the effect of parental income on children. *Journal of Political Economy*, 109(2):266–280, 2001.

[20] World Bank. *Edstats Database*. Accessed October, 2008.

A Appendix: Proof that Probability of Success Increases with δ and β

As shown by equation (4), when the chosen γ is at an interior solution where $0 \leq \gamma \leq 1$,

$$\gamma^* = \frac{b\delta^2 - a\beta^2}{2b\delta^2 - a\beta^2}$$

where I omit the time subscript for ease of exposition.

To prove that the probability of success is increasing in β or δ, it is sufficient to show that $\frac{(1-\gamma^*)\delta^2}{a}$ is also increasing in β or δ. This follows because on the interior, the first-order condition (3) is equivalent to

$$\frac{(1-\gamma^*)\delta^2}{a} = (\delta c_1 + \beta p_1) \tag{11}$$

To prove that $\frac{(1-\gamma^*)\delta^2}{a}$ is increasing in δ, differentiate with respect to δ :

$$-\frac{\partial\gamma^*}{\partial\delta}\frac{\delta^2}{a} + (1-\gamma^*)\frac{2\delta}{a} \tag{12}$$

where

$$\frac{\partial\gamma^*}{\partial\delta} = \frac{2ab\delta\beta^2}{\left(2b\delta^2 - a\beta^2\right)^2}$$

After substituting γ^* and $\frac{\partial\gamma^*}{\partial\delta}$ into (12) and simplifying, we have

$$\frac{\partial}{\partial\delta}\frac{(1-\gamma^*)\,\delta^2}{a} = \frac{4b^2\delta^5 - 4ab\beta^2\delta^3}{\left(2b\delta^2 - a\beta^2\right)^2}$$

Then, since on the interior we know that $\beta^2 < \delta^2\frac{b}{a}$, it follows that

$$4b^2\delta^5 - 4ab\beta^2\delta^3 > 0$$

since $4ab\beta^2\delta^3 < 4ab\left(\delta^2\frac{b}{a}\right)\delta^3$.

Because $\frac{(1-\gamma^*)\delta^2}{a}$ is decreasing in γ^*, we only need to show that γ^* is decreasing in β to show that the probability of success is increasing in β. We can do this by differentiating the solution for

γ^* with respect to β:

$$\frac{\partial \gamma^*}{\partial \beta} = -\frac{2a\beta}{2b\delta^2 - a\beta^2} - \frac{2a\beta\left(b\delta^2 - a\beta^2\right)}{\left(2b\delta^2 - a\beta^2\right)^2} < 0$$

Figure 1: Distributions of Raw Test Scores, Program and No-Program Groups

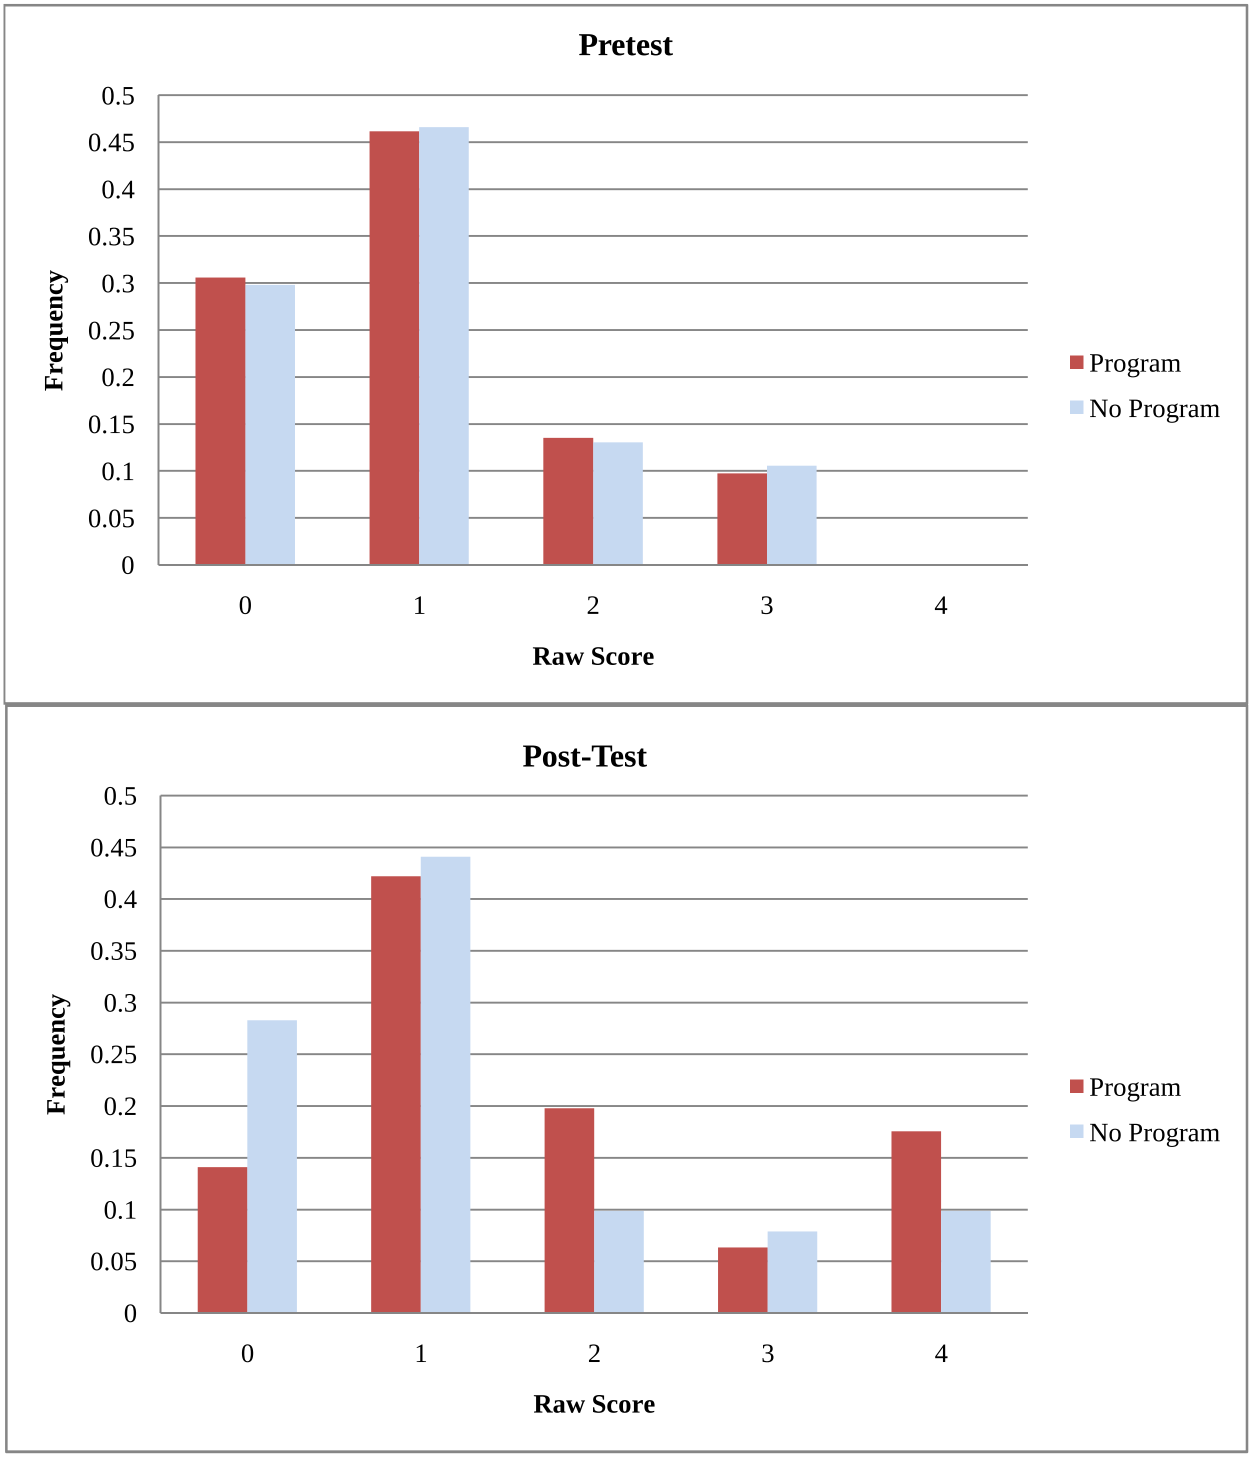

Figure 2: Treatment Effects of Toys by Pretest Score Attendance in After-school Classes

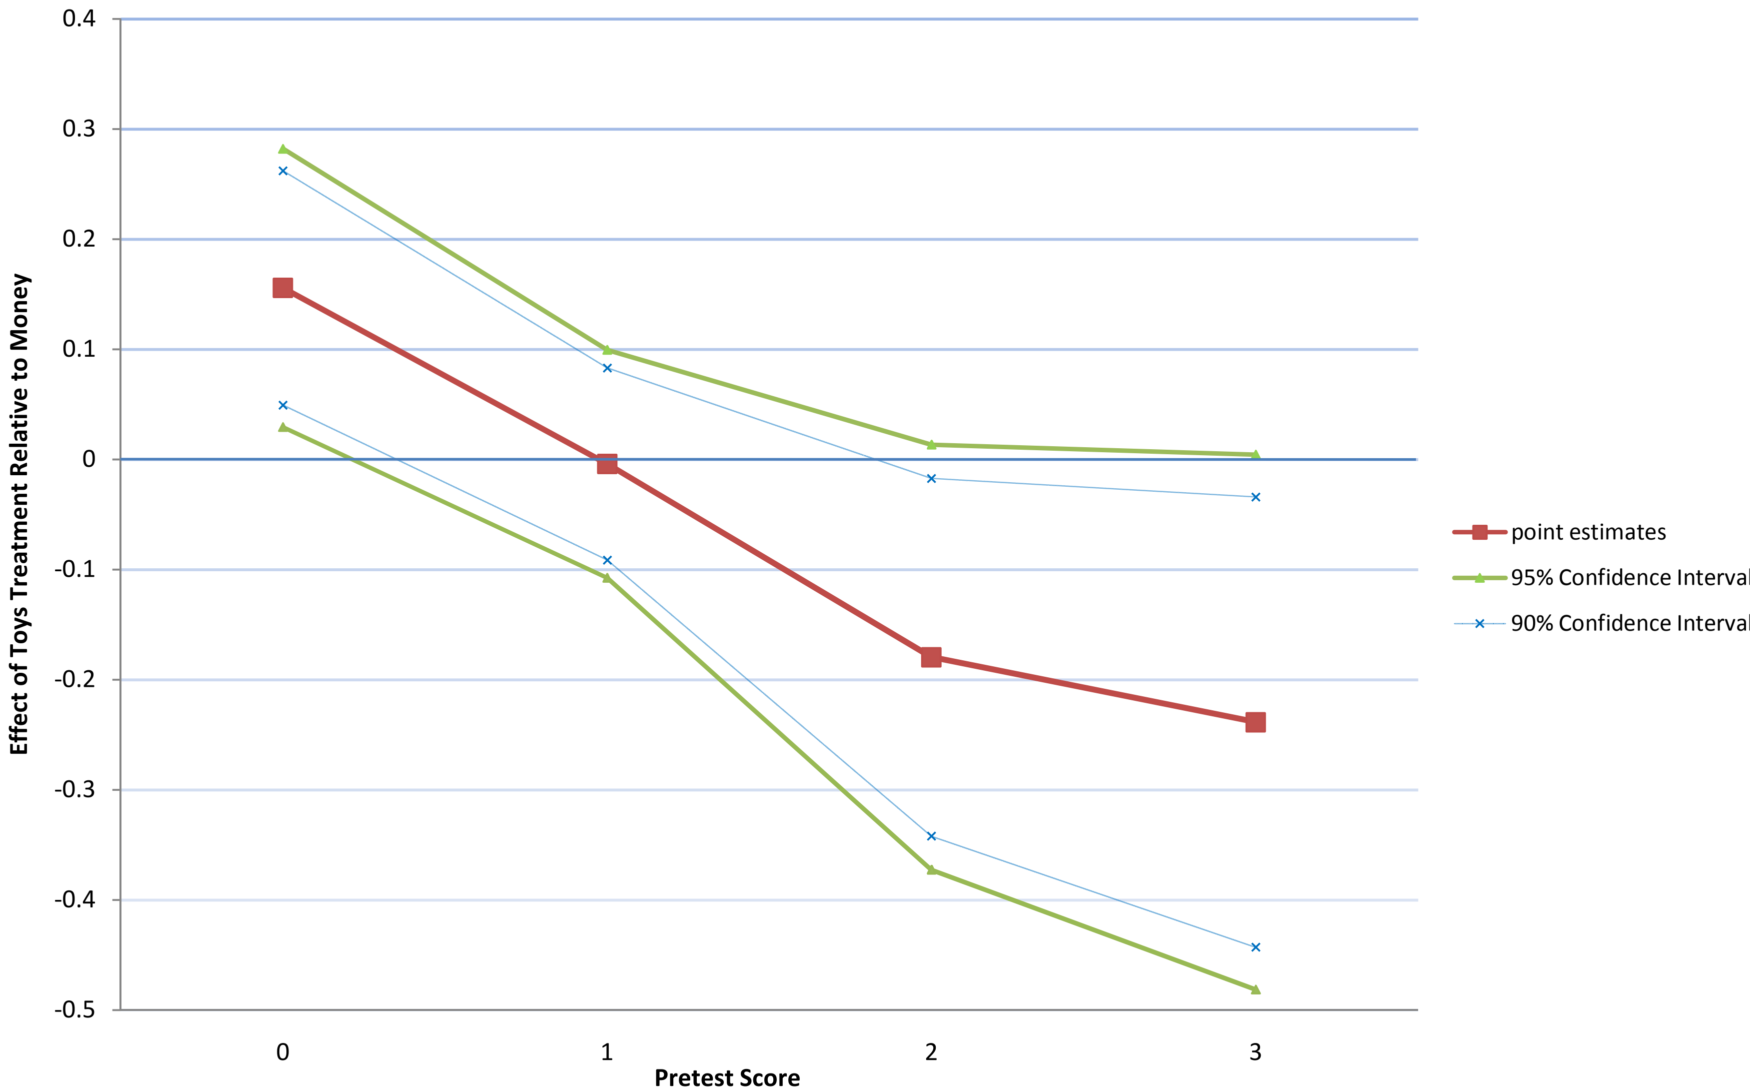

Note: This figure uses the coefficients from a regression of attendance on toys treatments interacted with pretest score, controlling for pretest score and classroom dummies.

Figure 3: Treatment Effects of Toys by Pretest Score Achievement of Literacy Goal

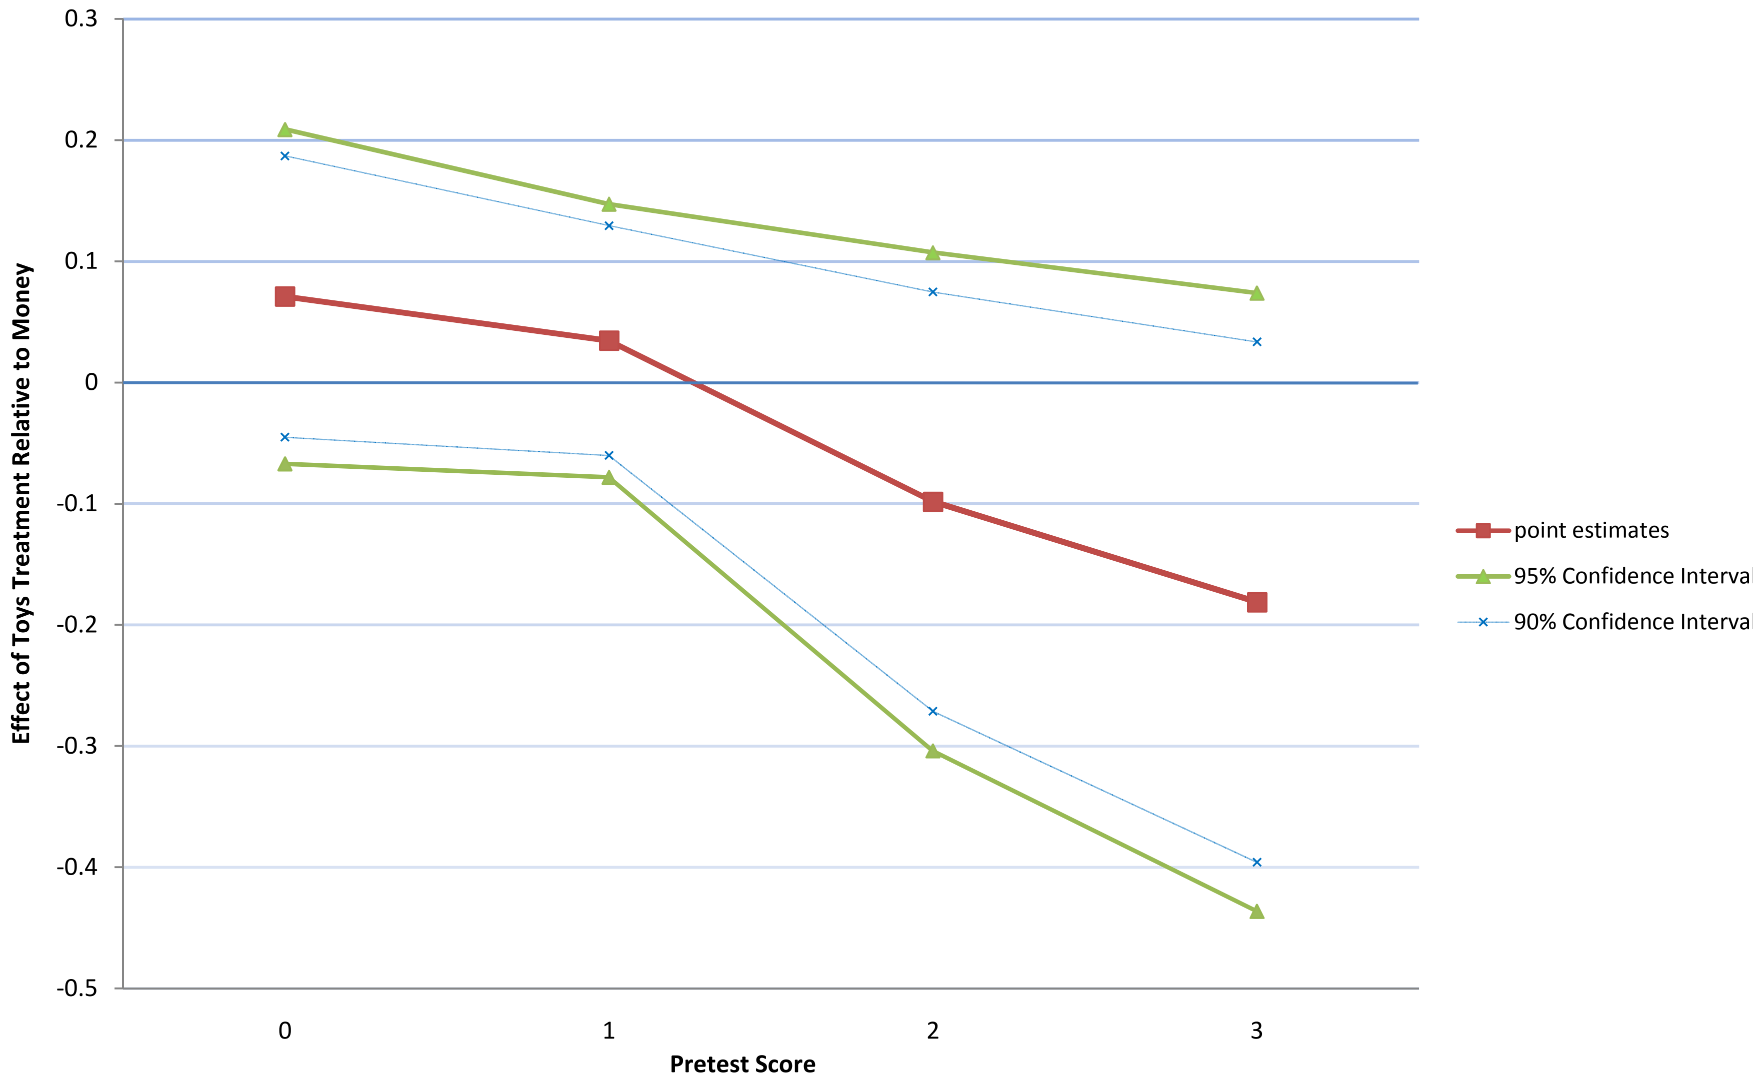

Note: This figure uses the coefficients from a regression of achievement on toys treatments interacted with pretest score, controlling for pretest score and classroom dummies.

Table 1: Sample Composition

	Reached at Pretest	Randomized Sample	Reached for Program Announcement	Reached at Post-test
	(1)	(2)	(3)	(4)
Total	1466	1086	925	900
Pretest Score				
0	349	331	283	276
1	528	502	427	414
2	151	146	125	124
3	107	107	90	86
4	331	0	0	0
Grade				
1	410	384	331	319
2	552	431	363	353
3	504	271	231	228
Treatment				
Parent Money		179	156	150
Child Money		181	156	152
Voucher		182	156	151
Child Toy		180	149	145
Ex Ante Choice		183	153	151
Ex Post Choice		181	155	151

Notes:
The randomized sample is slightly imbalanced across treatments because several schools were randomized individually, and the remainders were not balanced across these schools.

Table 2: Balance of Observables Across Treatment Groups

	Dependent Variable										
	Pretest Score	Female	Mother Education	Father Education	Durables	Total Household Members	Money Given	Gave Toys	Gave Other Item	Helped with Studies	Tutoring Fees Paid
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
<u>Panel A:</u> Full Sample											
Child Money	-0.071	-0.032	-0.28	0.12	0.003	-0.02	-0.564	-0.025	0.096+	0.007	-0.332
	(0.103)	(0.056)	(0.386)	(0.432)	(0.018)	(0.151)	(2.14)	(0.021)	(0.054)	(0.032)	(5.856)
Voucher	-0.103	0.006	0.226	0.595	-0.007	-0.141	-2.686	-0.025	0.077	-0.012	2.7
	(0.103)	(0.056)	(0.382)	(0.427)	(0.018)	(0.151)	(2.136)	(0.021)	(0.054)	(0.032)	(5.856)
Child Toy	-0.024	0.013	0.298	0.633	-0.007	-0.141	-3.117	-0.011	0.021	0.044	-2.692
	(0.105)	(0.057)	(0.386)	(0.430)	(0.018)	(0.152)	(2.161)	(0.022)	(0.055)	(0.032)	(5.935)
Ex Ante Choice	-0.038	0.011	-0.177	0.268	0.019	-0.047	-3.491	-0.045*	0.078	0.002	7.261
	(0.104)	(0.056)	(0.383)	(0.428)	(0.018)	(0.151)	(2.15)	(0.021)	(0.055)	(0.032)	(5.894)
Ex Post Choice	-0.006	0.004	0.122	0.573	-0.023	-0.053	-5.196*	-0.019	0.105+	0.026	-1.79
	(0.104)	(0.056)	(0.382)	(0.428)	(0.018)	(0.150)	(2.140)	(0.021)	(0.054)	(0.032)	(5.866)
Constant	1.064**	0.571**	3.144**	6.054**	0.268**	5.432**	15.577**	0.058**	0.295**	0.077**	25.719**
	(0.073)	(0.040)	(0.270)	(0.303)	(0.012)	(0.107)	(1.511)	(0.015)	(0.038)	(0.023)	(4.154)
Observations	925	925	907	889	913	892	924	923	925	922	915
R-squared	0.002	0.001	0.004	0.005	0.007	0.002	0.009	0.005	0.007	0.004	0.004
F	0.294	0.174	0.69	0.804	1.225	0.317	1.628	0.995	1.216	0.782	0.756
P-value	0.917	0.972	0.631	0.547	0.295	0.903	0.15	0.42	0.299	0.563	0.582
<u>Panel B:</u> Toys Treatments vs. Money Treatments											
Voucher + Toys	-0.029	0.026	0.399	0.555+	-0.009	-0.131	-2.615	-0.006	0.001	0.012	0.243
	(0.073)	(0.040)	(0.273)	(0.306)	(0.013)	(0.104)	(1.771)	(0.016)	(0.038)	(0.023)	(4.051)
Constant	1.029**	0.554**	3.007**	6.113**	0.270**	5.422**	15.295**	0.045**	0.343**	0.080**	25.552**
	(0.051)	(0.028)	(0.193)	(0.218)	(0.009)	(0.073)	(1.245)	(0.011)	(0.027)	(0.016)	-2.851
Observations	617	617	601	590	609	591	617	615	617	614	610
R-squared	0	0.001	0.004	0.006	0.001	0.003	0.004	0	0	0	0
F	0.156	0.418	2.129	3.28	0.461	1.595	2.18	0.116	0.001	0.281	0.004
P-value	0.693	0.518	0.145	0.071	0.497	0.207	0.14	0.733	0.973	0.597	0.952

Notes:
Each column represents a linear regression of the dependent variable on dummies for each treatment category.
In Panel A, the omitted category is parent money. In Panel B, the omitted categories are parent and child money.
+ significant at 10%; * significant at 5%; ** significant at 1%

Table 3: Means of Selected Demographics and Baseline Behaviors

	Variable	Mean
Demographics		
	Female	0.571
		(0.495)
	Mother's Education (years)	3.18
		(3.34)
	Father's Education (years)	6.42
		(3.69)
	Pct Durables Owned	0.266
		(0.154)
	Household Members 0-5	0.729
		(0.894)
	Household Members 6-14	2.19
		(0.94)
	Household Members over 14	2.44
		(0.82)
Behaviors Towards Program Child		
	Helped with studies (past day)	0.358
		(0.480)
	Money given (past week, in Rs)	13.08
		(18.90)
	Gave toys (past week)	0.037
		(0.188)
	Gave other item (past week)	0.088
		(0.283)
	Tutoring fees paid (past month)	26.58
		(51.35)

Notes:

Standard deviations in parentheses.

Durables include: bicycle, motorcycle, dvd player, radio, tv, refrigerator, gas stove, cooler, landline and mobile phone.

Table 4: Determinants of Being Reached At Baseline

	Dependent Variable: Reached At Baseline (Dummy)		
	(1)	(2)	(3)
Zpretest	-0.005	-0.017	-0.029*
	(0.015)	(0.015)	(0.013)
Grade 2	-0.007	-	-
	(0.026)		
Grade 3	-0.002	-	-
	(0.030)		
Female	0.028	0.094*	0.017
	(0.022)	(0.042)	(0.037)
Classroom Dummies	NO	YES	YES
Surveyor Dummies	NO	NO	YES
Observations	1052	1052	1052
R-squared	0.002	0.103	0.39

Notes:

"Zpretest" represents the difference between the child's score and the grade specific average, divided by the grade-specific standard deviation.

+ significant at 10%; * significant at 5%; ** significant at 1%

Table 5: Difference-in-Differences Estimates
Overall Effect of the Incentive Program

	Dependent Variable: Normalized Test Score		
	(1)	(2)	(3)
Any Incentive	-0.031	-0.024	-0.05
	(0.088)	(0.083)	(0.081)
Posttest	0.151	0.151	0.151
	(0.115)	(0.109)	(0.103)
Incentive*Posttest	0.511**	0.511**	0.511**
	(0.124)	(0.118)	(0.112)
Standard 2		-0.505**	
		(0.048)	
Standard 3		-0.808**	
		(0.054)	
Female		-0.142**	-0.116
		(0.042)	(0.078)
Classroom Dummies	NO	NO	YES
Surveyor Dummies	NO	NO	YES
Observations	2104	2104	2104
R-squared	0.091	0.187	0.284

Notes:
+ significant at 10%; * significant at 5%; ** significant at 1%

Table 6: Differences Between Treatments

	Dependent Variable: Attendance in After-School Class, 1/0			Dependent Variable: Achievement of Literacy Goal		
	(1)	(2)	(3)	(4)	(5)	(6)
Child Money	0.038	0.046	0.069	-0.04	-0.02	-0.016
	(0.053)	(0.052)	(0.053)	(0.058)	-0.055	-0.057
Toys	0.023	0.026	0.04	-0.013	-0.002	-0.005
	(0.046)	(0.045)	(0.046)	(0.050)	(0.048)	(0.050)
Pretest = 1		0.034	0.013		-0.253**	-0.262**
		(0.048)	(0.049)		(0.052)	(0.054)
Pretest = 2		-0.016	-0.051		-0.241**	-0.248**
		(0.069)	(0.072)		(0.072)	(0.076)
Pretest = 3		0.103	0.053		-0.092	-0.124
		(0.085)	(0.091)		(0.088)	(0.094)
Classroom Dummies	NO	YES	YES	NO	YES	YES
Controls	NO	NO	YES	NO	NO	YES
Observations	502	502	502	598	598	598
R-squared	0.001	0.155	0.239	0.001	0.190	0.248

Notes:

In the first three columns, the dependent variable is a dummy which equals 1 if the child attended the after-school class on at least 1 day. In Columns 4-6, the dependent variable is a dummy which equals 1 if the child reached the literacy goal.

The sample used in this table includes parent money, child money, voucher and toys treatments.

The "Toys" category includes both child toy and voucher groups. The omitted treatment category is parent money.

Controls include all variables in Table 2.

+ significant at 10%; * significant at 5%; ** significant at 1%

Table 7: Interactions of Toy Treatments and Pretest Scores
Outcome: Attendance in After-School Classes

	Using Actual Values of Zpretest			Zpretest Predicted by Parental Characteristics	
	(1)	(2)	(3)	(4)	(5)
Toys	-0.033	-0.034	-0.034	-0.134+	-0.160*
	(0.040)	(0.039)	(0.040)	(0.077)	(0.078)
Zpretest	-0.012	0.011	-0.014	0.187	0.258+
	(0.037)	(0.100)	(0.100)	(0.137)	(0.145)
Toys * Zpretest	-0.126**	-0.121**	-0.134**	-0.406*	-0.470*
	(0.046)	(0.045)	(0.046)	(0.202)	(0.202)
Pretest = 1	0.156**	0.095	0.118	0.098*	0.040
	(0.049)	(0.128)	(0.129)	(0.040)	(0.045)
Pretest = 2	0.148*	0.089	0.135	0.059	-0.008
	(0.071)	(0.212)	(0.214)	(0.060)	(0.069)
Pretest = 3	0.202*	0.225	0.247	0.103	0.091
	(0.081)	(0.270)	(0.272)	(0.074)	(0.094)
Classroom Dummies	NO	YES	YES	NO	YES
Addt'l Controls	NO	NO	YES	NO	NO
Observations	502	502	502	502	502
R-squared	0.037	0.167	0.204	0.03	0.192

Notes:
The dependent variable is a dummy which equals 1 if the child attended the after-school classes on at least one day.
"Zpretest" represents the difference between the child's score and the grade specific average, divided by the grade-specific standard deviation.
In Columns 4 and 5, Zpretest represents the predicted values of the regression in Column 1 of Appendix Table 1. Standard errors in these columns are constructed based on 500 bootstrap draws.
The sample used in this table includes parent money, child money, voucher and toys treatments.
The "Toys" category includes both toy and voucher groups. The omitted treatment categories are parent and child money.
+ significant at 10%; * significant at 5%; ** significant at 1%

Table 8: Interactions of Toy Treatments and Child's Initial Test Score
Outcome: Achievement of the Literacy Goal

	Using Actual Values of Zpretest			Zpretest Predicted by Parental Characteristics	
	(1)	(2)	(3)	(4)	(5)
Toys	-0.023	-0.02	-0.015	-0.060	-0.103
	(0.044)	(0.043)	(0.044)	(0.095)	(0.096)
Zpretest	-0.065	0.120	0.095	0.047	0.134
	(0.040)	(0.107)	(0.108)	(0.161)	(0.168)
Toys * Zpretest	-0.086+	-0.079	-0.060	-0.189	-0.304
	(0.051)	(0.050)	(0.050)	(0.243)	(0.238)
Pretest = 1	-0.057	-0.345**	-0.335**	-0.139**	-0.251**
	(0.053)	(0.129)	(0.129)	(0.049)	(0.053)
Pretest = 2	0.054	-0.399+	-0.374+	-0.080	-0.238**
	(0.077)	(0.215)	(0.217)	(0.065)	(0.074)
Pretest = 3	0.266**	-0.299	-0.272	0.122	-0.092
	(0.088)	(0.279)	(0.282)	(0.076)	(0.085)
Classroom Dummies	NO	YES	YES	NO	YES
Addt'l Controls	NO	NO	YES	NO	NO
Observations	598	598	598	598	598
R-squared	0.05	0.194	0.221	0.017	0.163

Notes:
The dependent variable is a dummy which equals 1 if the child achieved the literacy goal.
"Zpretest" represents the difference between the child's score and the grade specific average, divided by the grade-specific standard deviation.
In Columns 4 and 5, Zpretest represents the predicted values of the regression in Column 1 of Appendix Table 1. Standard errors in these columns are constructed based on 500 bootstrap draws.
The sample used in this table includes parent money, child money, voucher and toys treatments.
The "Toys" category includes both toy and voucher groups. The omitted treatment categories are parent and child money.
+ significant at 10%; * significant at 5%; ** significant at 1%

Table 9: Choice Between Toy and Money

Panel A: Choice Frequencies

	Sample		
	All	Achievers Only	
	Ex Ante Treatment	*Ex Ante* Treatment	*Ex Post* Treatment
	(1)	(2)	(3)
Chose Toy	0.327	0.316	0.506
Chose Money	0.673	0.684	0.494
Observations	153	76	79

Panel B: Determinants of Choice of Toy—*Ex Ante* Treatment

	Dependent Variable: Parent Chose Toy		
	(1)	(2)	(3)
Productivity Index	0.019	0.120	0.029
	(0.211)	(0.262)	(0.249)
Classroom Dummies	NO	YES	YES
Surveyor Dummies	NO	NO	YES
Observations	151	151	151
R-squared	0.0001	0.291	0.52

Notes:
The dependent variable in Panel B is a dummy which equals 1 if the parent chose the toy.
In Panel B, the Productivity Index is constructed as the predicted values of the regression in Column 1 of Appendix Table 1.
The sample used in Panel B includes children in the *ex ante* choice treatment.
+ significant at 10%; * significant at 5%; ** significant at 1%

Table 10: Interactions of Choices Treatments and Baseline Achievement

	Dependent Variable					
	Class Attendance, 1/0			Achieve Goal		
		Zpretest			Zpretest	
		Actual	Predicted		Actual	Predicted
	(1)	(2)	(3)	(4)	(5)	(6)
Toys	0	-0.035	-0.145+	0.014	-0.02	-0.066
	(0.037)	(0.039)	(0.076)	(0.040)	(0.043)	(0.094)
Ex Ante Choice	0.058	0.048	0.137	-0.024	-0.037	-0.015
	(0.045)	(0.047)	(0.099)	(0.049)	(0.052)	(0.120)
Ex Post Choice	-0.024	-0.03	-0.155	0.034	0.022	0.189
	(0.045)	(0.048)	(0.099)	(0.049)	(0.052)	(0.121)
Zpretest		0.075	0.212		0.082	0.118
		(0.083)	(0.138)		(0.088)	(0.160)
Toys * Zpretest		-0.116*	-0.423*		-0.085+	-0.212
		(0.046)	(0.199)		(0.049)	(0.233)
Ex Ante Choice * Zpretest		-0.022	0.232		-0.016	0.044
		(0.056)	(0.247)		(0.060)	(0.283)
Ex Post Choice * Zpretest		-0.011	-0.381		-0.033	0.461
		(0.054)	(0.267)		(0.058)	(0.305)
Classroom Dummies	YES	YES	YES	YES	YES	YES
Observations	755	755	755	900	900	900
R-squared	0.129	0.142	0.143	0.12	0.175	0.178

Notes:

"Zpretest" represents the difference between the child's score and the grade specific average, divided by the grade-specific standard deviation.

In Columns 2 and 4, the actual values of Zpretest are used.

In Columns 3 and 6, Zpretest represents the predicted values from the regression in Column 1of Appendix Table 1. Standard errors in these columns are constructed based on 500 bootstrap draws.

The "Toys" category includes both child toy and voucher groups. The omitted treatment categories are parent and child money.

+ significant at 10%; * significant at 5%; ** significant at 1%

Table 11: Difference-in-differences Estimates of Effect of Toys Treatments vs. Money Treatments

	Any Incentive	Money given	Gave Toys	Gave Other Item
	(1)	(2)	(3)	(4)
Panel A. First Follow-up vs. Baseline				
Toys	-0.035	-3.071*	-0.002	0.015
	(0.043)	(1.417)	(0.020)	(0.028)
Post	0.128**	-1.606	0.027	0.128**
	(0.043)	(1.393)	(0.019)	(0.027)
Toys*Post	-0.012	2.594	0.018	-0.080*
	(0.060)	(1.984)	(0.027)	(0.039)
Pretest Dummies	YES	YES	YES	YES
Classroom Dummies	YES	YES	YES	YES
Observations	1178	1192	1178	1176
R-squared	0.093	0.086	0.063	0.083
Panel B. Post-test Follow-up vs. Baseline				
Toys	-0.035	-3.278*	-0.004	0.016
	(0.048)	(1.521)	(0.024)	(0.031)
Post	0.520**	-0.08	0.135**	0.367**
	(0.049)	(1.576)	(0.023)	(0.031)
Toys*Post	-0.310**	2.297	-0.053	-0.275**
	(0.070)	(2.239)	(0.033)	(0.043)
Pretest Dummies	YES	YES	YES	YES
Classroom Dummies	YES	YES	YES	YES
Observations	1086	1095	1182	1180
R-squared	0.202	0.092	0.11	0.186

Notes:

The "Toys" category includes both child toy and voucher groups. The omitted treatment categories are parent and child money.

"Post" is a dummy variable for the later survey (either the first or post-test followup)

Variables were based on a one-week recall.

Money given was the total amount in Rs. given to the child over the past week.

The sample used in this table includes parent money, child money, voucher and toys treatments.

+ significant at 10%; * significant at 5%; ** significant at 1%

Appendix Table 1: Relationship Between Relative Pretest Score and Parental Productivity Measures

	Dependent Variable: Zpretest		
	(1)	(2)	(3)
# Children under 15	-0.051*	-0.039+	-0.038+
	(0.023)	(0.022)	(0.022)
# Adults 15+	0.017	0.059+	0.064*
	(0.033)	(0.031)	(0.031)
Pct durables owned	-0.613**	-0.383*	-0.391*
	(0.171)	(0.161)	(0.165)
Mother employed	-0.111*	-0.101+	-0.092+
	(0.055)	(0.052)	(0.054)
Mother education	0.016+	0.008	0.008
	(0.009)	(0.008)	(0.009)
Father education	-0.004	0	0
	(0.008)	(0.008)	(0.008)
Helped with studies	0.064	0.029	0.004
	(0.055)	(0.052)	(0.056)
Tutoring fees paid/10	0.018**	0.015**	0.015**
	(0.005)	(0.005)	(0.005)
School-Grade Dummies	NO	YES	YES
Surveyor Dummies	NO	NO	YES
Observations	925	925	925
R-squared	0.047	0.251	0.274

Notes:

"Zpretest" represents the difference between the child's score and the grade-specific average, divided by the grade-specific standard deviation.

+ significant at 10%; * significant at 5%; ** significant at 1%