

Incentives Work: Getting Teachers to Come to School*

Esther Duflo, Rema Hanna, and Stephen P. Ryan[†]

November 11, 2008

Abstract

We use a randomized experiment and a structural model to test whether monitoring and financial incentives can reduce teacher absence and increase learning in rural India. In treatment schools, teachers' attendance was strictly monitored using cameras, and their salaries were made a nonlinear function of attendance. In the treatment schools, absenteeism by teachers fell by 21 percentage points relative to the control group, and children's test scores increased by 0.17 standard deviations. A structural dynamic labor supply model estimated in the treatment group closely predicts behavior in the control group, and is used to compute cost-minimizing compensation policies.

*This project is a collaborative exercise involving many people. Foremost, we are deeply indebted to Seva Mandir, and especially to Neelima Khetan and Priyanka Singh, who made this evaluation possible. We thank Ritwik Sakar and Ashwin Vasan for their excellent work coordinating the fieldwork. Greg Fischer, Shehla Imran, Callie Scott, Konrad Menzel, and Kudzaisha Takavarasha provided superb research assistance. For their helpful comments, we thank Abhijit Banerjee, Rachel Glennerster, Michael Kremer and Sendhil Mullainathan. For financial support, we thank the John D. and Catherine T. MacArthur Foundation.

[†]The authors are from MIT (Department of Economics and J-PAL), the Kennedy School of Government of Harvard University and J-PAL, and MIT (Department of economics)

1 Introduction

Over the past decade, many developing countries have expanded primary school access. However, these improvements in school access have not been accompanied by improvements in school quality. For example, in India, a nationwide survey found that 65 percent of children enrolled in grades 2 through 5 in government primary schools could not read a simple paragraph (Pratham, 2006). These poor learning outcomes may be due, in part, to teacher absenteeism. Using unannounced visits to measure attendance, a nationally representative survey found that 24 percent of teachers in India were absent during normal school hours (Chaudhury, et al., 2005a, b).¹ Improving attendance rates may be the first step needed to make “universal primary education” a meaningful term.

Solving the absentee problem poses a significant challenge (see Banerjee and Duflo (2005) for a review). One solution—championed by many, including the 2004 World Development Report—is to involve the community in teacher oversight, including the decisions to hire and fire teachers. However, in many countries, teachers are a powerful political force, and may resist attempts to curb their influence. As such, many governments have begun to shift from hiring government teachers to instead hiring “para-teachers.” Para-teachers are teachers that are hired on short, flexible contracts to work in primary schools and in non-formal education centers (NFEs) run by non governmental organizations (NGOs) and local governments. Unlike government teachers, it may be feasible to implement greater oversight and incentives for para-teachers since they do not form a entrenched constituency, they are already subject to yearly renewal of their contract, and there is a long queue of qualified job applicants. Thus, providing para-teachers with incentives may be an effective way to improve the quality of education, provided that para-teachers can teach effectively.

In this paper, we use both experimental and structural methods to empirically test whether the direct monitoring of the attendance of para-teachers (referred to simply as teachers in the rest of the paper), coupled with high-powered incentives based on their attendance, improves school quality.

The effect of incentives based on presence is theoretically ambiguous. While simple labor supply models would predict that incentives should increase effort, the incentives could fail for a variety of reasons. First, the incentives may not be high enough. Second, the incentive schemes may crowd out the teacher’s intrinsic motivation to attend school (Benabou and Tirole, 2006). Finally, some teachers, who previously believed that they were required to

¹Although teachers do have some official non-teaching duties, this absence rate is too high to be fully explained by this particular story.

work every day, may decide to stop working once they have reached their target income for the month (Fehr and Gotte, 2002).

Even if incentives increase teacher attendance, it is unclear whether child learning levels will actually increase. Teachers may multitask (Holmstrom and Milgrom, 1991), reducing their efforts along other dimensions.² Such schemes may also demoralize teachers, resulting in less effort (Fehr and Schmidt, 2004), or may harm teachers' intrinsic motivation to teach (Kreps, 1997). On the other hand, incentives can improve learning levels if the main cost of working is the opportunity cost of attending school and, once in school, the marginal cost of teaching is low. In this case, an incentive system that directly rewards presence would stand a good chance of increasing child learning. Thus, whether or not the incentives can improve school quality is ultimately an empirical question.

We study a teacher incentive program run by the NGO Seva Mandir. Seva Mandir runs single-teacher NFEs in the rural villages of Rajasthan, India. Teacher absenteeism is high, despite the threat of dismissal for repeated absences. In our baseline study (August 2003), the absence rate was 44 percent. The program began in September 2003. In 57 randomly selected program schools, Seva Mandir gave teachers a camera, along with instructions to have one of the students take a picture of the teacher and the other students at the start and close of each school day. The cameras had tamper-proof date and time functions, allowing for the collection of precise data on teacher attendance that could be used to calculate teachers' salaries. Each teacher was then paid according to a nonlinear function of the number of valid school days for which they were actually present, where a "valid" day was defined as one for which the opening and closing photographs were separated by at least five hours and both photographs showed at least eight children. Specifically, they received Rs 500 if they attended fewer than 10 days in a given month, and Rs 50 for any additional day (up to a maximum of 25 or 26 days depending on the month). In the 56 comparison schools, teachers were paid a fixed rate for the month, and were told (as usual) that they could be dismissed for repeated, unexcused absences.

The program resulted in an immediate and long-lasting improvement in teacher attendance rates in treatment schools, as measured through monthly unannounced visits in both treatment and comparison schools. Over the 30 months in which attendance was tracked, teachers at program schools had an absence rate of 21 percent, compared to 44 percent

²This is a legitimate concern as other incentive programs (based on test scores) have been subject to multitasking (Glewwe, Ilias and Kremer, 2003), manipulation (e.g., Figlio and Winicki, 2002; Figlio and Getzler, 2002) or outright cheating (Jacob and Levitt, 2003). On the other hand, Lavy (2002) and Mulharidharan and Sundaraman (2006) find very positive effects of similar programs.

baseline and the 42 percent in the comparison schools.

While the reduced form results inform us that this program was effective in reducing absenteeism, it does not tell us what the effect of another scheme with a different payment structure would be. Moreover, it does not allow us to disentangle the effect of the monitoring from the effect of the incentives, since the comparison school teachers were not monitored. To answer these questions, we exploit the nonlinear nature of the incentive scheme to estimate a dynamic labor supply model using the daily attendance data in the treatment schools. The identification exploits the fact that the incentive for a teacher to attend school on a single day changes as a function of the number of days they attend school in the month, and the number of days left in the month. This is because they have to attend at least 10 days in a month to begin to receive the incentive (by working in the beginning of the month, the teacher builds up the option to work for Rs 50 per day at the end of the month). Indeed, regression discontinuity design estimates show that teachers work significantly more at the beginning of the month than at the end of the previous month, when they had not accumulated at least 10 days of work in that month.

We use this fact to estimate the teachers' marginal utility of money. We allow serial correlation in the opportunity cost of attending school and heterogeneity in teachers' outside option, and we use the method of simulated moments to estimate the parameters. Allowing for serial correlation and heterogeneity considerably complicates the estimation procedure, but we show that these features are very important in this application. To our knowledge, our paper is one of the few papers to estimate dynamic labor supply decisions with unobserved heterogeneity and a serially correlated error structure.³ Other papers combining randomized experiments and structural models include Wise (1985), Lise, Seitz and Smith (2004), Card and Hyslop (2008), Ferrall (2008), Todd and Wolpin (2006), and Attanazio, Meghir, and Santiago (2006).

We find that teachers are responsive to the financial incentives: our preferred estimates suggest that the elasticity of labor supply with respect to the level of the financial bonus is 0.306. Furthermore, *decreasing* the number of days that workers must work until they are eligible for the incentive by a single day *increases* the expected number of days worked by about 1.29 percent. We do not use the data from the control group in our structural model. This allows us to use the data to test whether the model accurately predicts teacher presence in the control group. The model that includes both serial correlation and teacher

³The only other paper we know of that allows for both of these factors is the working paper of Bound, Stinebrickner, and Waidman (2005), which examines the retirement decisions of older men.

heterogeneity does well in these out of sample tests: when we set the incentive to zero, it predicts the difference in attendance in the treatment and the control group, as well as the number of days worked under a new incentive system initiated by Seva Mandir after the experiment. We use the parameters of the model to compute the optimal incentive scheme.

Although we find that teachers are sensitive to the financial incentives, we see no evidence of multitasking. When the school was open, teachers were as likely to be teaching in treatment as in comparison schools, suggesting that the marginal costs of teaching are low conditional on attendance. Student attendance when the school was open was similar in both groups, so that the students in the treatment group received more days of instruction. A year into the program, test scores in the treatment schools were 0.17 standard deviations higher than in the comparison schools. Two and a half years into the program, children from the treatment schools were also 10 percentage points (or 62 percent) more likely to transfer to formal primary schools, which requires passing a competency test. The program's impact and cost are similar to other successful education programs.

The paper is organized as follows. Section 2 describes the program and evaluation strategy. The results on teacher attendance are presented in Section 3, while the estimates from the dynamic labor supply model are presented in Section 4. Section 5 presents the results on other dimensions of teacher effort, as well as student outcomes. Section 6 concludes.

2 Experimental Design and Data Collection

2.1 Non-formal Education Centers

Since the enactment of the National Policy on Education in 1986, non-formal education centers (NFEs) have played an important role in India's drive toward universal primary education. They have been the main instrument for expanding school access to children in remote and rural areas. They have also been used to transition children who may otherwise not attend school into a government school. As of 1997, 21 million children were enrolled in NFEs across India (Education for All Forum, 2000). Similar informal schools operate throughout most of the developing world (Bangladesh, Kenya, etc.).

Children of all ages may attend the NFE, though, in our sample, most are between 7 and 10 years of age. Nearly all of the children are illiterate when they enroll. In the setting of our study, the NFEs are open six hours a day and have about 20 students each. All students are taught in one classroom by one teacher, who is recruited from the local community and who has, on average, a 10th grade education. Instruction focuses on basic Hindi and math

skills. The schools only have one teacher; thus, when the teacher is absent, the school is closed.

2.2 The Incentive Program

Seva Mandir runs about 150 NFEs in the tribal villages of Udaipur, Rajasthan. Udaipur is a sparsely populated, hard to access region. Thus, it is difficult to regularly monitor the NFEs, and absenteeism is high. A 1995 study (Banerjee et al., 2005) found that the absence rate was 40 percent, while our baseline (in August 2003) found that the rate was 44 percent.

To reduce teacher absenteeism, Seva Mandir implemented an external monitoring and incentive program starting in September 2003. They chose 120 schools to participate, with 60 randomly selected schools serving as the treatment group and the remaining 60 as the comparison group.⁴ In the treatment schools, Seva Mandir gave each teacher a camera, along with instructions for one of the students to take a photograph of the teacher and the other students at the start and end of each school day. The cameras had a tamper-proof date and time function that made it possible to precisely track each school's openings and closings.⁵ Camera upkeep (replacing batteries, and changing and collecting the film) was conducted monthly at regularly scheduled teacher meetings. If a camera malfunctioned, teachers were instructed to call the program hotline within 48 hours. Someone was then dispatched to replace the camera, and teachers were credited for the missing day. Rolls were collected every two months, and payments were distributed every two months.

At the start of the program, Seva Mandir's monthly base salary for teachers was Rs 1000 (\$23 at the real exchange rate, or about \$160 at PPP) for at least 20 days of work per month. In the treatment schools, teachers received a Rs 50 (\$1.15) bonus for each additional day they attended in excess of the 20 days (where holidays and training days, or about 3 days per month on average, are automatically credited as working days), and they received a Rs 50 fine for each day of the 20 days they skipped work. Seva Mandir defined a "valid" day as one in which the opening and closing photographs were separated by at least five hours and at least eight children were present in both photos to indicate that the school was actually functioning. Due to ethical and political concerns, Seva Mandir capped the fine at Rs 500.

⁴After randomization but prior to the announcement of the program, 7 of these schools closed. The closures were equally distributed amount the treatment and controls schools, and were not due to the program. We thus have 57 treatment schools and 56 comparison schools.

⁵The time and date buttons on the cameras were covered with heavy tape, and each had a seal that would indicate if it had been tampered with. Fines would have been imposed if cameras had been tampered with (this did not happen) or if they had been used for another purpose (this happened in one case).

Thus, salaries ranged from Rs 500 to Rs 1,300 (or \$11.50 to \$29.50). In the 56 comparison schools, teachers were paid the flat rate of Rs 1,000, and were informed that they could be dismissed for poor attendance. However, this happens very rarely, and did not happen during the span of the evaluation.⁶

2.3 Data Collection

An independent evaluation team led by Vidhya Bhawan (a consortium of schools and teacher training institutes) and J-PAL collected the data. We have two sources of attendance data. First, we collected data on teacher attendance through one random unannounced visit per month in all schools. By comparing the absence rates obtained from the random checks across the two types of schools, we can determine the program's effect on absenteeism.⁷ Second, Seva Mandir provided us with access to the camera and payment data for the treatment schools.

We collected data on teacher and student activity during the random check. For schools that were open during the visit, the enumerator noted the school activities: how many children were sitting in the classroom, whether anything was written on the blackboard, and whether the teacher was talking to the children. While these are crude measures of teacher performance, they were chosen because each could be easily observed before the teachers could adjust their behavior. In addition, the enumerator also conducted a roll call and noted whether any of the absent children had left school or had enrolled in a government school, and then updated the evaluation roster to include new children.

To determine whether child learning increased as a result of the program, the evaluation team, in collaboration with Seva Mandir, administered three basic competency exams to all children enrolled in the NFEs in August 2003: a pre-test in August 2003, a mid-test in April 2004, and a post-test in September 2004. The pre-test followed Seva Mandir's usual testing protocol. Children were given either a written exam (for those who could write) or an oral exam (for those who could not). For the mid-test and post-test, all children were given both the oral exam and the written exam; those unable to write, of course, earned a zero on the written section. The oral exam tested simple math skills (counting, one-digit addition, simple division) and basic Hindi vocabulary skills, while the written exam tested

⁶Teachers in the control schools knew that the camera program was occurring, and that some teachers were randomly selected to be part of the pilot program.

⁷Teachers understood that the random checks were not linked with an incentive. We cannot rule out the fact that the random check could have increased attendance in comparison schools. However, we have no reason to believe that they would differentially affect the attendance of comparison and treatment teachers.

for these competencies plus more complex math skills (two-digit addition and subtraction, multiplication and division), the ability to construct sentences, and reading comprehension. Thus, the written exam tested both a child’s ability to write and his ability to handle material requiring higher levels of competency relative to the oral exam.

2.4 Baseline and Experiment Integrity

Given that schools were randomly allocated to the treatment and comparison groups, we expected school quality to be similar across groups prior to the program onset. Before the program was announced in August 2003, the evaluators were able to randomly visit 41 schools in the treatment group and 39 in the comparison.⁸ Panel A of Table 1 shows that the attendance rates were 66 percent and 64 percent, respectively. This difference is not statistically significant. Other measures of school quality were also similar prior to the program: in all dimensions shown in Table 1, the treatment schools appear to be slightly better than comparison schools, but the differences are always small and never significant. Finally, to determine the joint significance of the treatment variable on all of the outcomes listed in Panels B through D, we estimated a Seemingly Unrelated Regression (SUR) model. The F-statistic is 1.21, with a p-value of 0.27, implying that the comparison and treatment schools were similar to one another at the program’s inception.

Baseline academic achievement, as measured by the pre-test, was the same for students across the two types of schools (Table 1, Panel E). On average, students in both groups appeared to be at the same level of preparedness before the program. There is no significant difference in either probability to take the written test or scores on the written tests.

3 Results: Teacher attendance

3.1 Reduced form results: Teacher Behavior

The effect on teacher absence was both immediate and long lasting. Figure 1 shows the fraction of schools found open on the day of the random visit, by month. Between August and September 2003, teacher attendance increased in treatment schools relative to the comparison schools. Over the next two and a half years, the attendance rates in both types of

⁸Due to time constraints, only 80 randomly selected schools of the 113 were visited prior to the program. There was no significant (or perceivable) difference in the characteristics of the schools that were not observed before the program. Moreover, the conclusion of the paper remains unchanged when we restrict all the subsequent analysis to the 80 schools that could be observed before the program was started.

schools followed similar seasonal fluctuations, with treatment school attendance systematically higher than comparison school attendance.

As Figure 1 shows, the treatment effect remained strong even after the post-test, which marked the end of the formal evaluation. Since the program had been very effective, Seva Mandir maintained it. However, at the end of the study, they only had enough resources to keep the program operating in the treatment schools. The random checks conducted after the post-test showed that the higher attendance rates persisted at treatment schools even after the teachers knew that the program was permanent, suggesting that teachers did not alter their behavior simply for the duration of the evaluation.

Table 2 presents a detailed breakdown of the program effect on absence rates for different time periods. On average, the teacher absence rate was 21 percentage points lower (or about half) in the treatment than in the comparison schools (Panel A).⁹ The effects on teacher attendance were pervasive—teacher attendance increased for both low and high quality teachers. Panel B reports the impact for teachers with above median test scores on the teacher skills exam conducted prior to the program, while Panel C shows the impact for teachers with below median scores.¹⁰ The program impact on attendance was larger for below median teachers (a 24 percentage point increase versus a 15 percentage point increase). However, this was due to the fact that the program brought below median teachers to the same level of attendance as above median teachers (78 percent).

The program reduced absence everywhere in the distribution. Figure 2 plots the observed density of absence rates in the treatment and comparison schools for 25 random checks. The figure clearly shows that the program shifted the entire distribution of absence for treatment teachers.¹¹ Not one of the teachers in the comparison schools was present during all 25 observations. Almost 25 percent of teachers were absent more than half the time. In contrast, 5 of the treatment teachers were present on all days, 47 percent of teachers were present on 21 days or more, and all teachers were present at least half the time. Thus, the program was effective on two margins: it eliminated very delinquent behavior (less than 50 percent attendance), and increased the number of teachers with high attendance records.

⁹This reduction in school closures was comparable to that of a previous Seva Mandir program which tried to reduce school closures by hiring a second teacher for the NFEs. In that program, school closure only fell by 15 percentage points (Banerjee, Jacob and Kremer, 2005), both because individual teacher absenteeism remained high and because teachers did not coordinate to come on different days.

¹⁰Teacher test scores and teacher attendance are correlated: In the control group, below median teachers came to school 53 percent of the time, while above median teachers came to school 63 percent of the time.

¹¹We also graphed the estimated cumulative density function of the frequency of attendance, assuming that the distribution of absence follows a beta-binomial distribution (not shown for brevity). The results are similar to that of Figure 2.

A comparison of the random check data and the camera data suggests that, for the most part, teachers did not “game” the system. Out of the 1337 cases where we have both camera data and a random check for a day, 80 percent matched. In 13 percent of the cases, the school was found open during the random check, but the photos indicated that the day was not considered “valid”, often because the photos were not separated by five hours. There are 88 cases (7 percent) in which the school was closed and the photos were valid, but only 54 (4 percent of the total) of these were due to teachers being absent in the middle of the day during the random check and shown as present both before and after. In the other cases, the data did not match because the random check was completed after the school had closed for the day, or there were missing data on the time of the random check or photo.

One interesting question is whether the effect of the program would be very different in the long run, because the program would induce different teachers to join schools with cameras. As of October 2008, the program is still in place in the same schools (Seva Mandir will soon extend it to all schools). We monitored the schools during a year from September 2006 to September 2007. Several teachers were new at that time, and new teachers hired in cameras schools have somewhat higher test scores than those hired in non cameras schools. After four years, teacher attendance was still significantly higher in the camera schools (72 percent versus 61 percent). Thus, this program seems to have a very long lasting effect on teacher attendance.

3.2 The Impact of Financial Incentives: Preliminary Evidence

The program had two components: the daily monitoring of teacher attendance and an incentive that was linked to attendance. Ideally, to disentangle these effects, we would have provided different monitoring and incentive schemes in different, randomly selected schools. Some teachers could have been monitored, but without receiving incentives. Some could have received a small incentive, while others could have received a larger one. This was not feasible for a variety of reasons. However, the nonlinear nature of the incentive scheme provides us with an opportunity to try to isolate the effect of the financial incentives, if we are willing to assume that the effect of the threat of monitoring does not follow exactly the same time pattern. Consider a teacher who, because he was ill, was unable to attend school on most of the first 20 days of the 26 days of the month. By day 21, assuming he has attended only 5 days so far, he knows that, if he works every single day remaining in the month, he will have worked only 10 days. Thus, he will earn Rs 500, the same amount he would earn if he did not work any other days that month. Although he is still monitored (and may worry

that if he does not attend at all in a month he may be punished), his monetary incentive to work in these last few days is zero. At the start of the next month, the clock is re-set. He now has incentive to start attending school again, since by attending at the beginning of the month he can hope to be “in the money” by the end of the month, thereby benefiting from the incentive. Consider another teacher who has worked 10 days by the 21st day of the month. For every day he works in the five remaining days, he earns Rs 50. By the beginning of the next month, his incentive to work is no higher. In fact, it could even be somewhat lower since he may not benefit from the work done the first day of the month if he does not work at least 10 days in that month.

This leads to a simple regression discontinuity design test for whether financial incentives matter, under the assumption that the direct monitoring effect of the camera (or any shock to the outside option of the teachers) does not have the same monthly discontinuity. Figure 3 gives a graphical representation of the approach. It shows a regression of the probability that a teacher works if she is in the money by day 21 of the month (with 4 days left), in the last 10 days of that month and the first 10 days of the next month. We fit a third order polynomial on the left and the right of the change in month. The figure shows a jump up for teachers who were not in the money, and no jump for those who were in the money. This is exactly what we would expect: the change in incentive at the beginning of a month is important for teachers who were not in the money since, in the data, we see that 65 percent of the teachers who are out of the money in a month will be in the money the following month. The teachers who were in the money, however, have a 95 percent chance to be in the money again. In addition, these teachers value the fact that the first days worked help them work toward the 10 days threshold. Correspondingly, we do not see a sharp drop in presence for the teachers who had been in the money the previous month.

Table 3 presents these results in regression form. Specifically, for teachers in the treatment group, we created a dataset that contains their attendance records for the last and the first day of each month. The last day of each month and the next day of the following month form a pair, indexed by m . We run the following equation, where $Work_{itm}$ is a dummy variable equal to 1 if teacher i works in day t in the pair of days m (t is either 1 or 2):

$$Work_{itm} = \alpha + \beta 1_{im}(d > 10) + \gamma Firstday_t + \lambda 1_{im}(d > 10) * Firstday_t + v_i + \mu_m + \epsilon_{itm}, \quad (1)$$

where $1_{im}(d > 10)$ is a dummy equal to 1 for both days in the pair m if the teacher had worked more than 10 days in the month of the first day of the pair, and 0 otherwise. $Firstday_t$ is a dummy that indicates that this is the first day of month (i.e. the second day of the pair).

We estimate this equation treating v_i and μ_m as either fixed effects or random effects. If the teachers are sensitive to financial incentives, we expect β to be positive (teachers should work more when they are in the money than out of the money), γ to be positive (a teacher who is out of the money in a given month should work more in the first day of the following month) and λ to be negative and as large as γ (there is no increase in incentive for teachers who had worked at least 10 days before).¹²

The results clearly show that teachers are more likely (46 percentage points in the random effect specification in Column 3) to attend school at the beginning of a month if they were not in the money in the previous month, which we do not see for teachers who were in the money. This holds after controlling for teacher fixed effects (Column 4), and even if we restrict the sample to the first and last day of the month (Columns 1 and 2).

These results imply that teachers are responsive to the financial incentives, unless there are other factors affecting teachers that happen to have exactly the same structure. Shocks to teachers' outside options are unlikely to change discontinuously when the calendar month changes, as no other teacher activity is linked to the calendar month. The direct effect of being monitored with the camera (over and above the financial incentive) is also unlikely to vary with the day of the month. Seva Mandir does not make use of daily attendance records: the payment software just calculates the attendance for the month. They care about overall attendance, not about how it varies with the calendar day in the month. In fact, in the control group, we find no relationship between the calendar day in the month and the chance that we see a teacher at work.¹³ Of course, if teachers thought that Seva Mandir's probability to fire them changed discontinuously every month that their attendance was less than 10 days (and that 2 days was no worse than 10 days) it would be impossible to separate that from the incentives, since it would have exactly the same time-structure as the incentive. It is not likely to be the case, however. Seva Mandir's official policy is that teachers should attend at least 20 days per month (although they know that few teachers reach that performance). They do not consider that 10 days is a satisfactory performance: the fine is capped to provide insurance to the teacher in months when they have a bad shock.

¹²Note that even with teacher fixed effects, β does not have a causal interpretation, because shocks may be auto-correlated. For example, a teacher who has been sick the entire month, and thus has worked less than 10 days, may also be less likely to work the first day of the next month. However, because when a month starts and finishes is arbitrary and should not be related to the underlying structure of shocks, a positive γ indicates that teachers are sensitive to financial incentives, unless there is a common "first day of the month" effect unrelated to the incentives. A negative λ will be robust even to this effect, since it would suggest that only teachers who are "out of the money" experience a "first day of the month" effect.

¹³Result available from the authors.

Thus, these findings suggest that teachers do respond to incentives. However, without more structure, it is not possible to conclude whether the effect of the program is entirely due to incentives. To do so, we need to estimate the marginal utility of the incentive to the teacher, and therefore the marginal value of working on the first day of the month. To analyze this problem, we set up a dynamic labor supply model and we use the additional restrictions that the model provides to estimate its parameters.

4 A Dynamic Model of Labor Supply

We propose and estimate a simple partial equilibrium model of dynamic labor supply, which incorporates the teacher response to the varying incentives over a month.

Each month, there are T periods, where T is the number of days in the month.¹⁴

In each period, each teacher decides whether or not to attend school. Teachers are risk neutral and maximize expected utility. For the teachers in the treatment group the decision problem is dynamic, as the pay they receive at the end of the month is linked to how many days they worked in that month. Other than that, we assume that the game is the same in every period. We also assume that there is no discounting within the month.

There are three components to the model: the state space, the per-period payoff function, and the transitions between states. We denote the state space by s , where $s = (t, d)$, where $t = 1, \dots, T$ is the current time and d is the days worked previously in the current month. It follows that $d \leq t - 1$ in our notation. The transition process between states is trivial. At the end of each day, t increases by one, unless $t = T$, in which case it resets to $t = 1$. If a teacher has worked in that period d increases by one, otherwise it remains constant.

The payoff function consists of two parts: utility from not working and utility from the financial payoff at the end of the month. In each period the agent faces a binary decision to attend work or not. If the teacher does not go to work at time t , she receives a utility of:

$$\mu + \epsilon_t, \tag{2}$$

where μ is the average utility of not working and ϵ_t is a shock. If the teacher goes to work, she receives nothing.

¹⁴ T is 25 in most months. Note that out of the 25 days, there are also several days of training and holidays, which are automatically credited as days worked for the teachers by the payment algorithm. Our estimation procedure follows the same rule, but when we report the number of days worked, we report it out of the days where teachers actually had to make a decision, which is on average 22 days per month.

At the conclusion of the month T , the teachers in the treatment group also receive a financial payment that is a function of the number of days worked in the current month:

$$\pi(d) = 500 + \max\{0, d - 10\} \cdot 50. \quad (3)$$

Each teacher gets at least Rs 500, and for every day over 10 that a teacher works in the current month, he or she receives a bonus of Rs 50. The parameter β converts the financial payment into utility terms, so that the utility of the payment is $\beta\pi(d)$. Teachers in the control group receive $\pi(d) = 1000$ regardless of their attendance in that month.

Teachers in the treatment group are utility maximizers facing a dynamic programming problem, as working today does not give an immediate payoff but may result in higher utility at the end of the month. Given this payoff structure, for $t < T$, we can write the value function for each teacher as follows:

$$V(t, d) = \max\{\mu + \epsilon_t + EV(t + 1, d), EV(t + 1, d + 1)\}. \quad (4)$$

At time T , we have:

$$V(T, d) = \max\{\mu + \epsilon_T + \beta\pi(d) + EV(1, 0), \beta\pi(d + 1) + EV(1, 0)\}. \quad (5)$$

Note that the term $EV(1, 0)$ enters into both sides of the maximum operator in Equation 5. Since the expectation of this term is independent of any action taken today, without loss of generality we can ignore any dynamic considerations that arise in the next month when making decisions in the current month. This is useful since we can think about solving the value function by starting at time T and working backward, which breaks an infinite-horizon dynamic program into a repeated series of independent finite-time horizon dynamic programs. Teachers in the control group do not face dynamic incentives but much simpler repeated binary choice problem, where the model predicts they will go to work on days where $\mu + \epsilon < 0$.¹⁵

Equation 5 also motivates several normalizing assumptions. First, the choice between working and not working only depends on the difference in utilities; thus, we set the payoff to attending school to be equal to zero. Second, the mean of the shock and the mean level

¹⁵This embodies the assumption that the only reason why there are dynamic incentives in this model is because of the financial incentives. As discussed in the previous section, we rule out the possibility that Seva Mandir's monitoring intensity changes over time during the month or is itself a complicated nonlinear function of days worked during the month.

of utility of not working are not separately identified; as a result we set the mean of the shock to be equal to zero. Finally, Equation 5 is only identified up to scale, as multiplying both sides by a positive constant does not change the work decision. Therefore, we follow the discrete choice literature and set the variance of the error term equal to one.¹⁶

4.1 Estimators

We estimate several models of the general dynamic program described by Equations 4 and 5.¹⁷ The six models vary according to what we assume about the distribution of μ and ϵ . We start with the simplest i.i.d. model, and progressively add heterogeneity and auto-correlation. We estimate these models using only the treatment group data, because we do not have daily attendance data from the control group (we have just one observation per teacher and per month), and the estimation of our preferred model, which allows for auto-correlation of ϵ_t , requires daily data. However, following Todd and Wolpin (2006), we will use the control group as an out of sample check for the different results we find.

Models with i.i.d. Errors The simplest model that we estimate is one where all agents share the same marginal utility of income and average outside option of not working, and the shocks to the utility of not working are i.i.d. We use all of the days in the month in our estimation by utilizing the empirical counterpart of Equation 4 for $t < T$:

$$\begin{aligned} Pr(\text{work}; t, d, \theta) &= Pr(\mu + \epsilon_t + EV(t + 1, d) < EV(t + 1, d + 1)) \\ &= Pr(\epsilon_t < EV(t + 1, d + 1) - EV(t + 1, d) - \mu) \\ &= \Phi(EV(t + 1, d + 1) - EV(t + 1, d) - \mu), \end{aligned} \quad (6)$$

where $\Phi(\cdot)$ is the standard normal distribution. Each of the value functions in Equation 6 is computed using backward recursion from period T . The log-likelihood function for the model without serial correlation in the error terms is then:

$$LLH(\theta) = \sum_{i=1}^N \sum_{m=1}^{M_i} \sum_{t=1}^{T_m} [1(\text{work})Pr(\text{work}; t, d, \theta) + 1(\text{not work})(1 - Pr(\text{work}; t, d, \theta))], \quad (7)$$

where each agent is indexed by i , the months they work are indexed by $m = \{1, \dots, M_i\}$, and the days within each of those months are indexed by $t = \{1, \dots, T_m\}$. This likelihood is

¹⁶An alternative would have been to set $\beta = 1$ and to estimate the variance of the shock instead.

¹⁷We briefly discuss identification issues of these models below.

well-behaved and can be evaluated quickly since numerical integration is not necessary. This framework also allows us to relax the assumption that teachers all have the same μ , as it is natural to assume that different teachers face different average outside options. We estimate two models with unobserved heterogeneity, one with teacher fixed effects and one where the outside option is drawn from a normal distribution with unknown mean and variance.

Models with Serial Correlation It is reasonable to think that the shock to teacher’s outside option are not i.i.d. For example, when a teacher is sick, she may be sick for a few days.

Suppose that the shock follows an AR(1) process:

$$\epsilon_t = \rho\epsilon_{t-1} + \nu_t, \tag{8}$$

where ρ is the persistence parameter and ν_t is a draw from the standard normal distribution. Autocorrelation could be either positive (illness) or negative (teacher has a task to accomplish). Irrespective of whether ρ is positive or negative, we can no longer directly apply the estimator used in the i.i.d. case. The reason is that the ϵ_T will be correlated with d , as teachers with very high draws on ϵ_T are more likely to be in the region where $d < 10$ if ρ is positive (the converse will be true if ρ is negative). In this case, the expectation that $\epsilon = 0$ is invalid, and will bias our estimates of α and subsequently β .

Our solution to this problem uses the method of simulated moments. We integrate out over the unknown distribution of ϵ when solving the dynamic program. Our identifying assumption to address the “incidental parameters” problem is to treat the the draw on ϵ at $t = 1$ as coming from the limiting distribution of Equation 8, which is given by $\bar{F} = N(0, 1/(1 - \rho^2))$. Thus, the source of identification in our model is the discontinuity in payoff at the end of the month which is induced by the payment rule: the key identification assumption is that nothing else changes discontinuously at that time. In particular, as in the regression discontinuity figures we presented earlier, we assume that if there is a direct monitoring effect of the cameras, it does not change discontinuously at the end of the month.

We then solve the dynamic program defined by following:

$$V(t, d, \epsilon_t) = \max\{\mu + \epsilon_t + EV(t + 1, d), EV(t + 1, d + 1)\}. \tag{9}$$

At time T , we have:

$$V(T, d, \epsilon_T) = \max\{\mu + \epsilon_T + \beta\pi(d) + EV(1, 0), \beta\pi(d + 1) + EV(1, 0)\}, \quad (10)$$

The difference in these value functions as compared to the i.i.d case is that we now explicitly account for ϵ_t as a state variable, where the expectation of ϵ_{t+1} depends on ϵ_t through the persistence parameter ρ . Conditioning on the initial draw from \bar{F} , the econometrician can infer the distribution of ϵ conditional on where the teacher is in the state space.

To estimate a model with this stochastic process on the outside option, we discretize the error term into 200 states and solve the dynamic programming problem defined by Equations 9 and 10 for an initial guess of our parameters.¹⁸ We then simulate many work histories from this model, forming an unbiased estimate of the distribution of *sequences* of days worked at the beginning of each month. We simulate the model by drawing sequences $\epsilon = \{\epsilon_0, \dots, \epsilon_t\}$ and following the optimal policy prescribed the dynamic program. For different sequences of ϵ 's we obtain different work histories. Repeating this process many times results in unbiased estimates of the probabilities of all possible sequences. We then match the model's predicted set of probabilities over these sequences against their empirical counterparts. Denoting a sequence of days worked as A , we form a vector of moment conditions:

$$E[Pr(A; X) - \widehat{Pr}(A; X, \hat{\theta})] = 0, \quad (11)$$

where $Pr(A; X)$ is the empirical probability of observing a sequence of days worked conditioning on X , a vector containing the number of holidays and the maximum number of days in that month an agent could potentially work in that month.¹⁹ We form $\widehat{Pr}(A)$ through Monte Carlo simulation:

$$\widehat{Pr}(A; X, \hat{\theta}) = \frac{1}{M \cdot NS} \sum_{m=1}^M \sum_{i=1}^{NS} 1(A_i = A; X_m, \hat{\theta}), \quad (12)$$

where A_i is the simulated work history associated with simulation i , as derived from a dynamic program constructed in accordance with the parameters $\hat{\theta}$ and the characteristics

¹⁸The number of states for ϵ was determined by increasing the number of points in the discretization of the error term until there was no change in the expected distribution of outcomes. For alternative approaches to estimating dynamic discrete choice models with serially-correlated errors, see Keane and Wolpin (1994) and Stinebrickner (2000).

¹⁹This is necessary since the maximum payoff a teacher could obtain varies across months with the length of the month and the number of holidays in that month, which count as a day worked in the bonus payoff function if they fall on a workday.

of the month. The number of simulations used to form the expected probability of observing a sequence of days worked is denoted as NS . In all of our estimations we use $NS = 200,000$. Note that we are also drawing ϵ_1 anew from the distribution \bar{F} for every simulated path, where we keep track of the seeding values in the random number generator, as to ensure that the function value is always the same for a given $\hat{\theta}$.²⁰ The objective function under the method of simulated moments is:

$$\min_{\theta} \left[\sum_{i=1}^n g(X_i, \theta) \right]'^{-1} \Omega^{-1} \left[\sum_{i=1}^n g(X_i, \theta) \right], \quad (13)$$

where $g(X_i, \theta)$ is the vector of moments formed by stacking the $2^N - 1$ moments defined by Equation 11, and Ω^{-1} is the standard two-step optimal weighting matrix. For more details concerning the implementation and asymptotic theory of simulation estimators, see McFadden (1989) and Pakes and Pollard (1989).

Matching sequences of days worked from the first N days in each month produces $2^N - 1$ linearly independent moments, where we subtract one to correct for the fact that the probabilities have to sum to one. In our estimation, we match sequences of length $N = 5$, which generates 31 moments. Experimentation with shorter and longer sequences of days worked did not result in significant changes to the coefficients.²¹

We also relax the assumption that the outside option is equal across all agents. We model the distribution of mean levels of the outside option to each agent in each month, μ_{it} , as being drawn from an unknown distribution denoted by $G(\mu)$. When forming moments in the MSM estimator in Equation 13 we need to integrate out this unobserved heterogeneity.

²⁰The conditional distribution of ϵ in future months still depends on the actions taken today, the crux of the “incidental parameters” problem. However, the dependence across the first five days of one month to the first five days of the following month is quite weak, even with high values of ρ . In principle, we can test this restriction by estimating the model on just the first month’s worth of work sequences.

²¹There is also a related econometric problem: the more moments one has to match, the lower the number of observations corresponding to each sequence. As the number of moments gets large, the number of teachers who actually followed any specific sequence diminishes towards zero. The number of days we match reflects a tradeoff in the additional information embodied in a longer sequence of choice behavior against the empirical imprecision of measuring those moments. This is a conceptually separate problem from the computational burden of simulating the model probabilities precisely, which also contributes to noisy estimates. For example, using the first 16 days of the months, where we will start seeing some teachers “in the money,” generates 65,535 moments, which is more than double the number of observations in the data.

The modification to the expected probability of observing a sequence of actions, A_i , is then:

$$\widehat{Pr}(A; X, \widehat{\theta}) = \frac{1}{M \cdot NS \cdot U} \sum_{m=1}^M \sum_{i=1}^{NS} \sum_{u=1}^U 1(A_i = A; X_m, \widehat{\theta}_1, u), \quad (14)$$

where u is a draw of the mean level of the outside option from $G(\widehat{\theta}_2)$, the unknown distribution of heterogeneity in the population. In practice we set $U = 200$. For clarity, we have partitioned the set of unknown parameters into $\widehat{\theta}_1 = \{\beta, \rho\}$ and $\widehat{\theta}_2$, the set of parameters governing the distribution of unobserved heterogeneity. Note that this model is slightly different than the fixed-effects model considered in the i.i.d. case above, as it allows the draw of the outside option to vary across both months and agents. On the other hand, the distribution of unknown heterogeneity in the fixed effects model is estimated semiparametrically.

We estimate three different models with unobserved heterogeneity which differ through the specification of $G(\widehat{\theta}_2)$. In the first model $G(\widehat{\theta}_2)$ is distributed normally with mean and variance $\widehat{\theta} = \{\mu_1, \sigma_1^2\}$. In the second model, our preferred specification, we allow for a mixture of two types, where each type is distributed normally with proportion p and $(1 - p)$ in the population. The third model is identical to the second model except that we remove the autocorrelation in the error structure for comparison.

Note that the MSM estimator does not directly exploit the variation coming from the discontinuous change in incentives at the end of the month. However, the discontinuity and the nonlinear payment rule is still the source of identification in our model. As discussed above, there is no easy way to incorporate the discontinuity in the MSM framework since it would require using the work sequence over the entire month, which would generate too many moments to match.²²

4.2 Parameter Estimates

We present these results in Table 4. We present the main parameters of the model (β , μ , σ , and other parameters as relevant), as well as the implied labor supply elasticity, the percentage increase in the average number of days worked caused by a one percent increase

²²We do not claim this is the only way to estimate the model: there may be other estimation methods that directly exploit the discontinuity (one of our student is working precisely on this problem); however any alternative empirical approach, such as maximum simulated likelihood, has to confront the issue that choice behavior depends on the full history of previous choices, and the empirical probability of observing any sequence of long enough time is vanishingly small. Furthermore, the key identification assumption has to be the discontinuity at the end of the month, and our estimator produces consistent estimates under this assumption.

in the value of the bonus, and the semi-elasticity with respect to the bonus cut-off, the percentage increase in the average number of days worked in response to an increase in one day in the minimum number of days necessary for a bonus.

All of the models suggest that teachers place value on the financial incentives ($\beta > 0$). Note that all the models except for our preferred specification, Model V, which includes both serial correlation and heterogeneity, have positive estimates of μ . This suggests that these models are almost certainly miss-specified, as this implies teachers should not work on most days without financial incentives. However, we observe teachers to work more than half the time in the control group, suggesting that μ should be negative on average. In our preferred specification, which includes two types of heterogeneity and serial correlation, most teachers have outside options drawn from a distribution with negative mean ($\mu = -0.43$). A small minority of teachers ($p = 0.02$) have opportunity costs drawn from a very high positive μ , reflecting that proportion of the population which does not attend school very often (recall that approximately two percent of our sample does not work any days in any given month). Relatedly, all models except model V have high β and implausibly high semi-elasticities with respect to the bonus threshold.

The parameters estimates of model V explain why the models that do not incorporate both heterogeneity and serial correlation overestimate the β and μ . First, we estimate a high positive auto-correlation in the shock to the outside option (0.448). The three i.i.d. models (I, II and VI) attribute all persistence in days worked to the response to the incentives, which leads to an overestimate of β , since some of it is due to the persistence in the shock itself.²³ Second, there is important teacher heterogeneity. In model V we allow the distribution of outside options in the data to be distributed as a mixture of two normals. Thus, we can empirically investigate the hypothesis that there are “slackers” as well as hard-working teachers. The results of Model V strongly suggest that there are (at least) two types of workers in the data. The average level of the outside option for the first group is $\mu_1 = -0.428$, with a variance of $\sigma_1^2 = 0.007$. We estimate that a majority of teachers, ($1 - p = 0.98$), derive some utility of going to school in most months, as they have a negative outside option in most periods. On the other hand, we estimate that two percent of the population ($p = 0.02$) draws an outside option from a distribution with much higher mean $\mu_2 = 1.78$ and variance $\sigma_2^2 = 0.0497$. By ignoring this heterogeneity, models III and IV, which allow for serial correlation but assumed either no heterogeneity (III) or that the teachers drew their outside

²³Note that by ignoring auto-correlation in the error term, Model VI also misses the heterogeneity between teachers.

option every month from a single normal distribution (IV), overstated the μ for the majority of teachers. As a consequence, the other models also estimate a very high response to the incentives through β , which then leads to a large under-prediction of how many days the control group should work.

In Model V, the elasticity with respect to the bonus is 0.196 and the semi-elasticity with respect to the bonus cutoff is -0.14 percent, both estimated highly significantly.²⁴ Both of these numbers are relatively small; the reason is that with two types, most teachers either have relatively high or low outside options, and are therefore somewhat insensitive to marginal incentives. A second contributing factor is the relatively low β estimated in this model.

4.3 Policy Rules Simulations and Out of Sample Tests

To provide a sense of the fit of each model, we report the predicted number of days worked under each specification. This is not a good test for the models estimated using maximum likelihood (Models I, II, and VI), which use all the days worked to compute the parameters of the model, and should therefore do a good job of matching the average number of days worked. However, note that this is not a parameter that our method of simulated moments estimation tried to match (since we matched only the first five days of teacher behavior), so it provides a partial goodness of fit metric for these models. The actual average days worked, not counting holidays, was 17.16. All of the models match this figure very well, including our preferred specification (16.28 days per month). All of the predictions are precise as well, with the exception of Model VI, which was estimated with more noise and therefore has correspondingly less precise predictions across the board.²⁵

Figure 4A plots the density of days worked predicted by Model V, and its 95 percent confidence interval, and compares it to the actual density observed in the data. Since the estimation is not calibrated to match this shape (we are only using the history of the first few days in our estimation), the fit is surprisingly good. The model reproduces the general shape of the distribution, although the mode of the distribution in our predicted fit is to the left of the mode in the data by one day. The model tends to slightly over-predict the frequency of 17 to 21 days worked and under-predict the frequency of days worked between

²⁴The standard errors on all predictions were formed by simulating the models for 50 draws of the parameters from their asymptotic distributions.

²⁵The model estimates and standard errors are slightly different than the working paper version, as we were able to harness additional computational resources to increase the number of simulations in computing the models and their predictions. The reduction in standard errors reflects a decrease in the simulation error.

3 and 10. With the exception of a small proportion of teachers who work a small number of days in a month, the true distribution lies comfortably within the 95 percent confidence interval of the prediction.

A unique feature of this experiment is that we have two real out-of-sample tests to assess the fit of the model. The first test is to compute what would happen if the financial incentives were set to zero, as in the control group. The estimates vary significantly from one model to the next, and we have already seen evidence that most models were not good fits for the data, so we focus on our preferred specification (Model V). This model suggests that the teachers would work on average 12.9 out of 22 days if there were no incentives, or 58.6 percent. In the random check data, we found that control group teachers are present 58 percent of the time (Table 2).²⁶ In contrast to the other models, Model V therefore does an excellent job of predicting behavior in the control group. If our identification assumption that the daily pattern of attendance is due to the changing incentives during the month for an income-maximizing teacher is correct, this suggests that financial incentives entirely explain the effect of the program.

A change in the incentive system at Seva Mandir, after the first version of this paper was written and our model was estimated, provides us with a second very nice counterfactual experiment. In December 2006, Seva Mandir increased the minimum monthly payment to Rs 700, which teachers receive if they work 12 days or less (rather than 10 days). For each additional day they work, teachers earn an additional Rs 70 per day. Seva Mandir provided us with the camera data in the summer 2007, a few months after the change in policy. The average number of days worked since January 2007 increased very slightly, from 17.16 to 17.39 days. The predicted number of days worked for each model is reported in the last row in Table 4. Here again, our preferred specification (Model V) performs very well: it predicts 17.77 days worked under the new incentive scheme, an increase from the predicted number of days under the main scheme (as in the actual data). Figure 4B shows the actual distribution of days worked and the predicted one. The model does a good job of predicting the distribution of days worked in the out-of-sample test, although the empirical distribution has more variance than in the original experiment. In contrast, the other models continue to perform rather poorly, again due to the fact that the implied elasticities in these models are too high.

²⁶The model predicts 76 percent presence in the treatment group, and treatment teachers are present 79 percent of the time in random checks

4.4 Counterfactual Optimal Policies

A primary benefit of estimating a structural model of behavior is the ability to calculate outcomes under economic environments not observed in the data. In our case, we are interested in finding the cost-minimizing combination of the two policy instruments, the size of the bonus and the threshold to get into the bonus, that lead to a minimum number of days worked in a month. Using Model V as our foundation, we calculated the expected number of days worked and expected size of the financial payout for a wide range of potential policies under our preferred model with autocorrelation and two types of heterogeneity. We let the minimum number of days to obtain a bonus range from zero to 23, which is the upper limit of days that a teacher could work in any month. At the same time, we varied the bonus paid for each day over the cutoff from zero to 300 Rp/day in increments of 25 Rp/day. Table 5 shows the lowest-cost combinations of those two policy variables that achieved a minimum expected number of days worked.²⁷

The results of this simulation show two general trends: the cost-minimizing cutoff generally decreases and the bonus increases in the expected number of days worked that the policymaker wants to achieve. Both of these trends lead to drastically increasing costs as the target increases. This result directly follows from the model: as we continue to increase the target, the marginal teacher has increasingly higher opportunity costs of working. This becomes quite expensive, as soon it is necessary to incentivize the “slacker” teacher types in our sample. It is interesting to note that for about the same amount of money as spent on both the treatment and control groups in the actual experiment (roughly 1000 Rp/month), teachers under the optimal counterfactual policy would have worked approximately 20 days, an improvement of roughly 16 percent and 56 percent over the treatment group and control group, respectively. Our counterfactual calculations show that while the actual intervention was successful in increasing teacher attendance, the NGO could have induced higher work effort with approximately the same amount of expenditures by doubling the bonus threshold and nearly tripling the per-day bonus. This is due to the fact that teachers in our sample appear to be on average more likely than not to attend school even without incentives and are forward looking. A higher threshold avoids rewarding infra-marginal days, and provides incentives to teachers to work to accumulate the number of days necessary to get the larger prize.

²⁷The expected outcomes were subject to a very small amount of variance as we drew model primitives from their estimated distributions 50 times for each combination of policy instruments.

5 Was Learning Affected?

5.1 Teacher Behavior

Though the program increased teacher attendance and the length of the school day, it could still be considered ineffective if the teachers compensated for increased attendance by teaching less. We used the activity data that was collected at the time of the random check to determine what the teachers were doing once they were in the classroom. Since we can only measure the impact of the program on teacher performance for schools that were open, the fact that treatment schools were open more may introduce selection bias. That is, if teachers with high outside options (who are thus more likely to be absent) also tended to teach less when present, the treatment effect may be biased downward since more observations would be drawn from among low-effort teachers in the treatment group than in the comparison group. Nevertheless, Table 6 shows that there was no significant difference in teacher activities: across both types of schools, teachers were as likely to be in the classroom, to have used the blackboard, and to be addressing students when the enumerator arrived. This does not appear to have changed during the duration of the program.

The fact that teachers did not reduce their effort in school suggests that the fears of multitasking and loss of intrinsic motivation were perhaps unfounded. Instead, our findings suggest that once teachers were forced to attend (and therefore to forgo the additional earnings from working elsewhere or their leisure time), the marginal cost of teaching must have been small. This belief was supported during in-depth conversations with 15 randomly selected NFE teachers regarding their teaching habits in November and December 2005. We found that teachers spent little time preparing for class. Teaching in the NFE follows an established routine, with the teacher conducting the same type of literacy and numeracy activities every day. One teacher stated that he decides on the activities of the day as he is walking to school in the morning. Other teachers stated that, once they left the NFE, they were occupied with household and field duties, and, thus, had little time to prepare for class outside of mandatory training meetings. Furthermore, despite the poor attendance rates, many teachers displayed a motivation to teach. They stated that they felt good when the students learned, and liked the fact they were helping to educate disadvantaged students.

The teachers' general acceptance of the incentive system may be an additional reason why multitasking was not a problem. Several months into the program, teachers filled out feedback forms. Seva Mandir also conducted a feedback session at their bi-annual sessions, which were attended by members of the research team. Overall, teachers did not complain

about the principle of the program, although many teachers had some specific complaints about the inflexibility of the rules. For example, many did not like the fact that a day was not valid even if a teacher was present 4 hours and 55 minutes (the normal school day is six hours, but an hour’s slack was given). On the other hand, many felt empowered by the fact that the onus of performing better was actually in their hands: “Our payments have increased, so my interest in running the center has gone up.” Others described how the payment system had made other community members less likely to burden the teacher with other responsibilities once they knew that a teacher would be penalized if he did not attend school. This suggests that the program may actually have stronger effects in the long run, as it signals a change in the norms of what teachers are expected to do.

5.2 Child Presence

On the feedback forms, many teachers claimed that the program increased child attendance: “This program has instilled a sense of discipline among us as well as the students. Since we come on time, the students have to come on time as well.” Unfortunately, conditional on whether a school was open, the effect of the program on child attendance cannot be directly estimated without bias, because we can only measure child attendance when the school is open. For example, if schools that were typically open also attracted more children, and the program induced the “worst” school (with fewer children attending regularly) to be open more often in the treatment schools than in the comparison schools, then this selection bias will tend to bias the effect of the program on child attendance downwards. The selection bias could also be positive, for example if the good schools generally attract students with better earning opportunities, who are more likely to be absent, and the “marginal” day is due to weak schools catering to students with little outside opportunities. Selection bias is a realistic concern (and likely to be negative) since, for the comparison schools, there is a positive correlation between the number of times a school is found open and the number of children found in school. Moreover, we found that the effect of the program was higher for schools with originally weak teachers, which may attract fewer children.

Keeping this caveat in mind, child attendance was not significantly different in treatment and comparison schools. In Table 7, we present the child attendance rates in an open school, by treatment status (Panel A). An average child’s attendance rate was the same in treatment and comparison schools (46 percent). Excluding children who left the NFE, child attendance is higher overall (62 percent for treatment and 58 percent for comparison schools), but the difference is not significant.

However, treatment schools had more teaching days. Even if the program did not increase child attendance on a particular day, the increase in the number of days that the school was open should result in more days of instruction per child. The program's impact on child instruction time is reported in Panel B of Table 7. Taking into account days in which the schools were closed, a child in a treatment school received 9 percentage points (or 30 percent) more days of instruction than a child in a comparison school. This corresponds to 2.7 more days of instruction time a month at treatment schools. Since there are roughly 20 children per classroom, this figure translates into 54 more child-days of instruction per month in the treatment schools than in comparison schools. This effect is larger than that of successful interventions that have been shown to increase child attendance (Glewwe and Kremer, 2005; Banerjee, Jacob and Kremer, 2005). The effect on presence does not appear to be affected by student ability (proxied by the whether or not the child could take a written test in the pre-test). While presence increased slightly more for those who could not write prior to the program (14 versus 10 percentage points), this difference is not significant.

In summary, since children were as likely to attend class on a given day in the treatment schools as in the comparison schools, and because the school was open more often, children received significantly more days of instruction in the treatment schools. This finding suggests that the high teacher absence rate we observed is not likely to be the efficient response to a lack of interest by the children: if it were the case that children came to school 55 percent of the time because they could not afford to attend more than a certain number of days, then we would see a sharp reduction in child attendance in treatment schools on days when the school was open. On the other hand, we do not see a sharp increase in the attendance of children in the treatment schools. This suggests that either the teacher absence rate is not the main cause of the children's irregular attendance, or that the children have not yet had time to adjust. The latter explanation is not entirely plausible, however, since the program has now been in place for over two years, and we do not see a larger increase in the attendance of children in the later periods than in the earlier period.

5.3 Child Learning

Children in the treatment schools, on average, received about 30 percent more instruction time than children in the comparison schools, with no apparent decline in teacher effort. Some, however, argue that because para-teachers are less qualified than other teachers, it is not clear that they are effective despite the support and in-service training they get from NGOs like Seva Mandir. If para-teachers are indeed ineffective, the fact that it is possible to

induce them to attend school more often is not particularly relevant for policy. Understanding the effect of the program on learning is, therefore, critical.

5.3.1 Attrition and Means of Mid- and Post-Test

Before comparing test scores in the treatment and comparison schools, we must first ensure that selective attrition does not invalidate the comparison. There are two possible sources of attrition.²⁸ First, some children leave the NFEs, either because they drop out of school altogether or because they start attending regular primary schools. Second, some children were absent on testing days. To minimize the impact of attrition on the study, we made considerable attempts to track down the children (even if they had left the NFE to attend a formal school or had been absent on the testing day) and administered the post-test to them. Consequently, attrition was fairly limited. Of the 2,230 students who took the pre-test, 1,893 also took the mid-test, and 1,760 also took the post-test. Table 8 shows the attrition rate in both types of schools, as well as the characteristics of the attriters. At the time of the mid-test, attrition was higher in the comparison group than in the treatment group. At the time of the post-test, attrition was similar across both groups, and children who dropped out of the treatment schools were similar in terms of test scores to those who dropped out of the comparison schools.

Table 8 also provides some simple descriptive statistics, comparing the test scores of treatment and comparison children. The first row presents the percentage of children who were able to take the written exam, while subsequent rows provide the mean exam score (normalized by the mid-test comparison group). Relative to the pre-test and mid-test, many more children, in both the treatment and comparison schools, were able to write by the post-test. On the post-test, students did slightly worse in math relative to the mid-test comparison, but they performed much better in language.

Finally, Table 8 also shows the simple differences in the mid- and the post-test scores for students in the treatment and comparison schools. On both tests, in both language and math, the treatment students did better than the comparison students (a 0.16 standard deviation increase and 0.11 standard deviations in language at the post-test score), even though the differences are not significant. Since child test scores are strongly auto-correlated, we obtain greater precision by controlling for the child's pre-test score level.

²⁸As mentioned earlier, seven centers closed down prior to the start of the program. We made no attempt to test the children from these centers in the pre-test.

5.3.2 Test Results

In Table 9, we report the program’s impact on test scores. We compare the average test scores of students in the treatment and comparison schools, conditional on a child’s pre-program competency level. In a regression framework, we model the effect of being in a school j that is being treated (Treat_j) on child i ’s score (Score_{ijk}) on test k (where k denotes either the mid- or post-test exam):

$$\text{Score}_{ijk} = \beta_1 + \beta_2 \text{Treat}_j + \beta_3 \text{Pre_Writ}_{ij} + \beta_4 \text{Oral_Score}_{ij} + \beta_5 \text{Written_Score}_{ij} + \varepsilon_{ijk}. \quad (15)$$

Because test scores are highly autocorrelated, controlling for a child’s test scores before the program increases the precision of our estimate. However, the specific structure of the pre-test (i.e., there is not one “score” on a comparable scale for each child because the children either took the written or the oral test in the pre-test) does not allow for a traditional difference-in-difference (DD) or “value added” (child fixed effect) strategy. Instead, we include a variable containing the child’s pre-test score for the oral test if he took the oral pre-test and 0 otherwise (Oral_Score_{ij}), the child’s pre-test score on the written test if he took the written test and 0 otherwise ($\text{Written_Score}_{ij}$), and an indicator variable for whether he took the written test at the pre-test (Pre_Writ_{ij}).²⁹ This fully controls for the child pre-test achievement, and is thus similar in spirit to a DD strategy. Standard errors are clustered by school. Each cell in Table 9 represents the treatment effect (β_2) obtained in a separate regression. For ease of interpretation, the mid-test results (Columns 1 to 4) and post-test results (Columns 5 to 8) are expressed in the standard deviation of the distribution of the mid-test score in the comparison schools.³⁰

The tables reveal that the program had a significant impact on learning, even as early as the mid-test. Children in treatment schools gained 0.16 standard deviations of the test score distribution in language, 0.15 standard deviations in math, and 0.17 overall (Panel A). Including controls for school characteristics—location, teacher test scores, and the infrastructure index of school—does not significantly change our findings (Panel B). Children

²⁹At the pre-test, children were given either the oral or the written score. At the mid- and post-test, every child took the oral part, and every child who could write took the written exam (all children were given a chance to try the written exam; if they could not read, they were given a zero for the written test).

³⁰Scores are normalized such that the mean and standard deviation of the comparison group at the time of the mid-test exam is zero and one, respectively. (Specifically, we subtract the mean of the comparison group in the pre-test, and divide by the standard deviation.) This allows for comparison across samples, as well as with the results from other studies. We could not normalize with respect to the pre-test score distribution since not every child took the same test at the pre-test.

who could write at the time of the pre-test gained the most from the program. For example, they had mid-line test scores 0.25 standard deviations higher in treatment schools than in comparison schools (Panel D). Interestingly, the children who could write at the time of the pre-test do not increase their attendance rate in response to the greater teacher attendance rate relatively more than those who could not write at the time of the pre-test. Therefore, it is not that they have relatively more days of schooling than the students who could not write as a result of the program, but rather that they seem better equipped to make the most out of the additional days of schooling that they receive.

We compare the program's impact on girls versus boys in Panels E and F. Girls gained as much, if not more, from the program as boys. On the mid-test, 7 percentage-points more of girls in the treatment schools were able to write relative to the comparison schools, compared to only 2 percentage-points of boys (this five percentage point difference is significant).

The differences between students in the treatment and comparison schools persisted in the post-test (Columns 5 to 8). Children in treatment schools gained 0.21 standard deviations in language, 0.16 in math, and 0.17 overall (Panel A). Similar to the mid-test, much of the gains came from children who could write at the time of the pre-test. The post-test also suggests that girls gained slightly more from the program than the boys, but these differences are not significant. The treatment effect of 0.17 standard deviations is similar to other successful educational interventions, such as the Balsakhi Remedial Education Program in India during its first year (Banerjee, et al., 2005).

Finally, we examined other sources of heterogeneity in program effect (school infrastructure, teacher level of education, teacher proficiency measured by the test scores). The results (omitted for brevity, but available from the authors), suggest little or no heterogeneity in program effect along these dimensions.

5.3.3 Leaving the NFE

NFEs prepare children, who might not otherwise attend school, to enter government schools at the age-appropriate grade level. To do so, children must demonstrate proficiency for a grade, either by passing an exam or through vetting by a government teacher. The ability of his students to join government schools is, therefore, a strong signal of success for a NFE teacher. The program increased the graduation rate to the government schools. As shown in Table 10, 26 percent of students in the treatment schools graduated to the government schools, compared to only 16 percent in the comparison schools (by February 2006). This 10 percentage point difference implies a 62 percent increase in the graduation rate and is

significant.

In the final row of Table 10, we present the dropout rates for children who left school entirely (i.e. left the NFE and did not join a government school). The dropout rate is slightly lower for the treatment schools, but this difference is insignificant.

5.3.4 Estimating the Effect of Teacher Presence on Learning

The previous sections presented the reduced form analysis of the effect of the incentive program on child learning. Table 11 interprets what these estimates can tell us about the impact of teacher attendance.³¹ Columns 1 to 3 report simple correlations between the teacher attendance rate and the child test scores. Specifically, we report the coefficient estimate of the number of times a school was found open ($Open_j$) on a regression of either the mid-test or post-test scores:

$$Score_{ijk} = \beta_1 + \beta_2 Open_j + \beta_3 Pre_Writ_{ij} + \beta_4 Oral_Score_{ij} + \beta_5 Written_Score_{ij} + \varepsilon_{ijk}. \quad (16)$$

We continue to control for the child’s pre-test score and to cluster standard errors by school.

Column 1 reports OLS estimation of Equation 2 for the comparison schools. In this case, the random check data are used to estimate the number of times a school is found open. The coefficient is 0.20, indicating that the test scores of children in centers open 100 percent of the time would be 0.10 standard deviations higher than those of children in a center open 50 percent of the time. Note that this coefficient is insignificant.

This point estimate is similar to those reported in other studies (Chaudhury, et al., 2005a) and, taken at face value, would imply that the effect of teacher attendance on learning is not that large. Chaudhury et al. (2005a) conjectures that the measurement of absence rates based on a few random visits per school have considerable error, and may thus bias the results downwards. Consistent with this theory, the effect on the post-test scores, where having more months of random check data allows us to better estimate the absence rate per school, becomes larger (0.58 standard deviations). Our study provides a much more direct test of this hypothesis, since, for treatment teachers, the photograph data gives us the actual attendance. We present the OLS estimate of the effect of attendance for treatment teachers using the random check data (Column 2) and camera data (Column 3). Overall, the effect of teacher attendance is larger in the treatment schools than the comparison schools (0.39 in Column 2 to 0.20 in Column 1, both obtained with random check data). More interestingly,

³¹This estimate are the effect of being present at a random check, which combines the effect of having come at all, and having come for a longer time.

consistent with the measurement error hypothesis, the effect of teacher attendance is larger and much more significant when using the more accurate measure of attendance from the camera data, especially for the mid-test scores (the estimate is 0.87 standard deviations in the Column 3 as compared to 0.39 in Column 2). For the post-test, where we have a much more accurate measure of attendance from the random check data, the results from the two methods are similar (0.98 in Column 3 versus 1.17 in Column 2).

Finally, in Column 4, we pool both samples and instrument Open_j (as measured by the random check) with the treatment status of the school to obtain exogenous variation in the percentage of time the school was found open. Since we have shown that the program had a direct effect on the length of the school day, as well as whether or not the school opened at all, the 2SLS estimate captures the joint effect of outright absence and of a longer school day. The 2SLS estimates are higher than the OLS results found in Column 1, and they are indistinguishable from the OLS results in Column 3, obtained with the precisely measured absence rate. This suggests that the relatively low correlation between teacher absence and test scores that was observed in previous studies is indeed likely to be due to measurement error in the teacher absence data. The more precise IV estimates suggest that even a 10 percentage point reduction in the absence rate would result in a 0.10 standard deviation increase in child test scores.

Extrapolating these estimates (which must be done with caution, since the local effect may be different from the overall effect), we can conclude that the effect of being enrolled in an NFE for a year with a teacher present every day is about one standard deviation. This point estimate is similar to the effect of attending remedial education classes with a para-teacher for one year in urban India for children who are enrolled in regular primary school, but have not yet achieved basic numeracy or literacy (Banerjee et al., 2005). In that study, the point estimate was 1.12 standard deviations. Both of these studies therefore suggest that para-teachers can be extremely effective teachers, at least when an NGO provides them with proper training.

6 Conclusion

In this paper, we show that direct monitoring, combined with simple and credible financial incentives based on teacher attendance, leads to large increases in attendance among para-teachers in informal schools. Absenteeism fell from an average of 42 percent in the comparison schools to 21 percent in the treatment schools, without affecting the teachers' effort while

in school. As a result, the students in treatment schools benefited from about 30 percent more instruction time. The program had an economically significant impact on test scores: after one year, child test scores in program schools were 0.17 standard deviations higher than in comparison schools. Children were also much more likely to be admitted to government schools.

This paper contributes to a small but growing literature that exploits structural modelling and a carefully controlled randomized experiments to answer an economic question. On the substantive front, our results suggest that providing incentives for attendance in non-formal schools can increase learning levels. However, the question arises as to whether incentive programs can be instituted for government teachers. Teachers in government schools are often more politically powerful than para-teachers. Thus, it may prove difficult to institute a system in which they would be monitored daily using a camera or similar device. However, our findings suggest that the barriers currently preventing teachers from attending school regularly (e.g., distance, other activities, lack of interest by children) are not insurmountable. Given political will, it is possible that solutions to the absence problem could be found in government schools as well.³²

A recent experiment demonstrates the external validity of these results outside the NGO context, as well as the difficulty of extending such programs in government settings (Banerjee, Duflo and Glennerster, 2007). Following the results of the cameras program, the government of Rajasthan created a similar system for government nurses, whose absence rate was about 44 percent. The nurses were monitored using time and date stamps. The announced incentive system was severe: it called for a 50 percent reduction in the pay of nurses who were absent 50 percent of the time, and termination of persistently absent nurses. In the first few months, when these punishments were carried out, the program led to about a 50 percent reduction in absenteeism. However, after a few months, the government started granting a large number of ex-post “exemptions” (although the monitoring did continue). The absence rate in the treatment group quickly converged to that of the control group. This further confirms that monitoring is effective, but only when coupled with real incentives, as is suggested by the results of our structural model.

The program for nurses suggests that barriers exist to the implementation of incentive systems for government employees. However, our findings also imply that para-teachers can be effective. If implementing monitoring within the government system turns out to be

³²A sign that these results did indeed generate interest in the government is that Seva Mandir was awarded the annual Government of India “Digital Learning Award” of 2007 for this project.

impossible, our results provide support for the policy of increasing teaching staff through the hiring of para-teachers.

References

- Attanasio, Orazio, Costas Meghir and Ana Santiago (2006). "Education Choices in Mexico. Using a Structural Model and a Randomized Experiment to Evaluate Progresca," Mimeo, UCL.
- Banerjee, Abhijit, Esther Duflo and Rachel Glennerster (2007). "Putting Band-Aid on a Corpse. Incentives for nurses in the Indian Public Health Care System" *Journal of the European Economic Association* Vol. 6(2-3): 487-500.
- Banerjee, Abhijit, and Esther Duflo (2006). "Addressing Absence," *Journal of Economic Perspectives* 20(1): 117-132.
- Banerjee, Abhijit, Suraj Jacob and Michael Kremer, with Jenny Lanjouw and Peter Lanjouw (2005). "Moving to Universal Education: Costs and Tradeoffs" Mimeo, MIT.
- Banerjee, Abhijit, Shawn Cole, Esther Duflo and Leigh Linden (2007). "Remedying Education: Evidence from Two Randomized Experiments in India," *Quarterly Journal of Economics*, 122(3): 1235-1264, August.
- Benabou, Roland, and Jean Tirole (2006). "Incentives and pro-social Behavior" *American Economic Review*, 95(6): 1652-1678
- Bound, John, Todd Stinebrickner, and Timothy A. Waidman (2005). "Using a Structural Retirement Model to Simulate the Effect of Changes to the OASDI and Medicare Programs," ERIU Working Paper, University of Michigan.
- Card, David and Dean Hyslop (2005.) "Estimating the Effects of a Time-Limited Earnings Subsidy for Welfare Leavers" *Econometrica*, 73(6): 1723-1770.
- Chaudhury, Nazmul, Jeffrey Hammer, Michael Kremer, Karthik Muralidharan, F. Halsey Rogers (2005). "Teacher Absence in India: A Snapshot," *Journal of the European Economic Association*.3(2): 658-667.
- Education for All Forum (2000). *EFA Country Assessment Country Reports*.


- Ferral , Christopher (2008). “Explaining and Forecasting Results of the Self-Sufficiency Project” Queen’s Economics Department Working Paper No. 1165.
- Glewwe, Paul and Michael Kremer (2005). ”Schools, Teachers, and Education Outcomes in Developing Countries,” , forthcoming in *Handbook on the Economics of Education*.
- Fehr, Ernst, and Schmidt (2004). “Fairness and Incentives in a Multi-task Principal–Agent Model” *Scandinavian Journal of Economics*, 106(3), 453–474.
- Holmstrom Bengt and P. Milgrom (1991). “Multi-Task Principal-Agent Problems: Incentive Contracts, Asset Ownership and Job Design,” *Journal of Law, Economics and Organization*, VII: 24–52.
- Keane, Michael and Ken Wolpin (1994). “The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence,” *Review of Economics and Statistics*, 76: 648-672.
- Kreps, David (1997). “Intrinsic Motivation and Extrinsic Incentives,” *American Economic Review*, 87(2): 359-364.
- Lise , Jeremy, Shannon Seitz and Jeffrey Smith (2004). “Equilibrium Policy Experiments nad the evaluation of social Programs” NBER working paper no. 10283
- McFadden, D. (1989). “A Method of Simulated Moments for Estimation of Multinomial Discrete Response Models Without Numerical Integration,” *Econometrica*, 57:995-1026.
- Pakes, A. and D. Pollard, (1989). “Simulation and the Asymptotics of Optimization Estimators,” *Econometrica*, 57: 1027-1057.
- Pratham (2006) *Annual Status of Education Report, 2005*
- Stinebrickner, Todd (2000). “Serially Correlated State Variables in Dynamic, Discrete Choice Models,” *Journal of Applied Econometrics*, 15, 595-624.
- Todd, Petra and Kenneth Wolpin (2006). “Using Experimental Data to Validate a Dynamic Behavioral Model of Child Schooling and Fertility: Assessing the Impact of a School Subsidy Program in Mexico:,” *American Economic Review*, 2006, 96(5): 1384–1417.

Table 1: Baseline Data

	Treatment (1)	Control (2)	Difference (3)
<i>A. Teacher Attendance</i>			
School Open	0.66	0.64	0.02 (0.11)
	41	39	80
<i>B. Student Participation (Random Check)</i>			
Number of Students Present	17.71	15.92	1.78 (2.31)
	27	25	52
<i>C. Teacher Qualifications</i>			
Teacher Test Scores	34.99	33.62	1.37 (2.01)
	53	56	109
<i>D. Teacher Performance Measures (Random Check)</i>			
Percentage of Children Sitting Within Classroom	0.83	0.84	0.00 (0.09)
	27	25	52
Percent of Teachers Interacting with Students	0.78	0.72	0.06 (0.12)
	27	25	52
Blackboards Utilized	0.85	0.89	-0.04 (0.11)
	20	19	39
Fstat(1,110)			1.21
p-value			(0.27)
<i>E. Baseline Test Scores</i>			
Took Written Exam	0.17	0.19	-0.02 (0.04)
	1136	1094	2230
Total Score on Oral Exam	-0.08	0.00	-0.08 (0.07)
	940	888	1828
Total Score on Written Exam	0.16	0.00	0.16 (0.19)
	196	206	402

Notes: (1) Teacher Performance Measures from Random Checks only includes schools that were open during the random check. (2) Children who could write were given a written exam. Children who could not write were given an oral exam. (3) Standard errors are clustered by school.

Figure 1: Percentage of Schools Open during Random Checks



Note: (1) The program began in September 2003. August only includes the 80 schools checked before announcement of program. September includes all random checks between August 25 through the end of September. (2) Child learning levels were assessed in a mid-test (April 2004) and a post-test (November 2004). After the post-test, the "official" evaluation period ended. Random checks continued in both the treatment and control schools.

Table 2: Teacher Attendance

Sept 2003-Feb 2006			Difference Between Treatment and Control Schools		
Treatment	Control	Diff	Until Mid-Test	Mid to Post Test	After Post Test
(1)	(2)	(3)	(4)	(5)	(6)
<i>A. All Teachers</i>					
0.79	0.58	0.21	0.20	0.20	0.23
		(0.03)	(0.04)	(0.04)	(0.04)
1575	1496	3071	882	660	1529
<i>B. Teachers with Above Median Test Scores</i>					
0.78	0.63	0.15	0.15	0.15	0.14
		(0.04)	(0.05)	(0.05)	(0.06)
843	702	1545	423	327	795
<i>C. Teachers with Below Median Test Scores</i>					
0.78	0.53	0.24	0.21	0.14	0.32
		(0.04)	(0.05)	(0.06)	(0.06)
625	757	1382	412	300	670

Notes: (1) Child learning levels were assessed in a mid-test (April 2004) and a post-test (November 2004). After the post-test, the "official" evaluation period was ended. Random checks continued in both the treatment and control schools. (2) Standard errors are clustered by school. (3) Panels B and C only include the 109 schools where teacher tests were available.

**Figure 2: Impact of the Cameras
(out of at least 25 visits)**

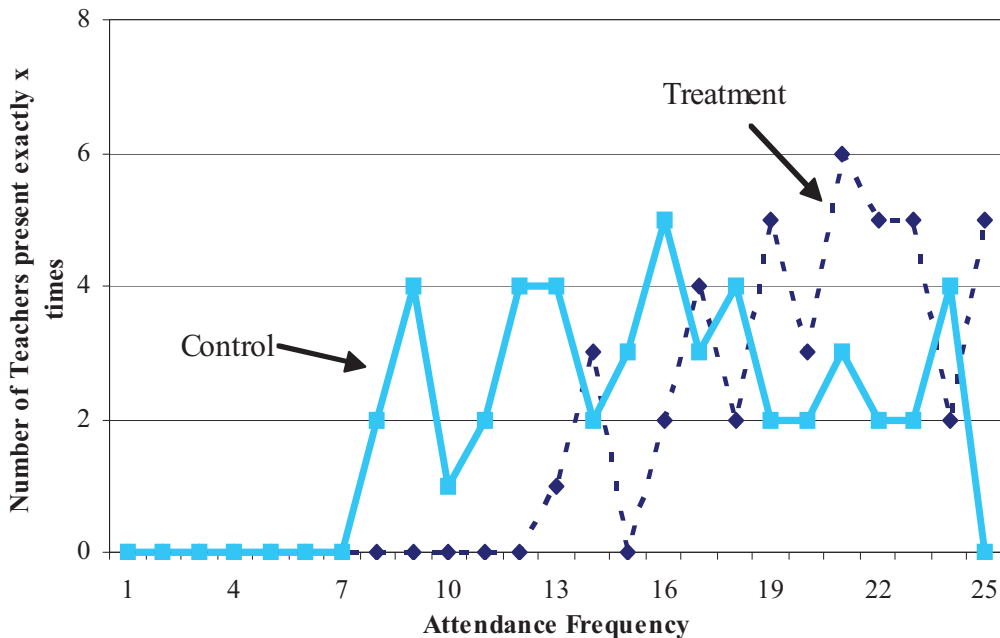



Table 3 : Do Teachers Work More When They are "In the Money"?

	(1)	(2)	(3)	(4)
Beginning of Month	0.19 (0.05)	0.12 (0.06)	0.46 (0.04)	0.39 (0.03)
In the Money	0.52 (0.04)	0.37 (0.05)	0.6 (0.03)	0.48 (0.01)
Beginning of the Month * In the Money	-0.19 (0.06)	-0.12 (0.06)	-0.34 (0.04)	-0.3 (0.02)
Observations	2813	2813	27501	27501
R-squared	0.06	0.22	0.08	0.16
Sample	1st and last day of month	1st and last day of month	1st 10 and last 10 days of month	1st 10 and last 10 days of month
Third Order Polynomial on Days on each side			X	X
Teacher Fixed Effects		X		X
Month Fixed Effects		X		X
Clustered Standard Errors	X		X	

Note: (1) The dependent variable in all models is an indicator variable for whether the teacher worked on a particular day, as measured by the photographs for the treatment schools.

Figure 3: RDD Representation of Teacher Attendance at the Start and End of the Month



Note: (1) The blue lines represent the months in which the teacher is in the money, while the red line represents the months in which the teacher not in the money. (2) The estimation includes a third order polynomial of days on the left and right side of the change of month.

Table 4: Results from the Structural Model

Parameter	Model I (1)	Model II (2)	Model III (3)	Model IV (4)	Model V (5)	Model VI (6)
β	0.049 (0.001)	0.027 (0.000)	0.055 (0.001)	0.057 (0.000)	0.013 (0.001)	0.043 (0.002)
μ_1	1.564 (0.013)		1.777 (0.013)	1.778 (0.021)	-0.428 (0.045)	1.236 (0.094)
ρ			0.422 (0.030)	0.412 (0.021)	0.449 (0.043)	
σ_1^2				0.043 (0.012)	0.007 (0.019)	0.143 (0.039)
μ_2					1.781 (0.345)	1.870 (0.672)
σ_2^2					0.050 (0.545)	2.051 (0.305)
p					0.024 (0.007)	0.003 (0.002)
Heterogeneity	None	FE	None	RC	RC	RC
ϵ_{Bonus}	1.09 (0.147)	0.592 (0.062)	1.299 (0.123)	1.82 (0.136)	0.196 (0.053)	1.371 (0.436)
$\epsilon_{\text{bonus_cutoff}}$	-18.26 (2.023)	-1.90 (0.564)	-16.94 (0.889)	-14.07 (1.609)	-0.14 (0.144)	-6.06 (3.635)
Predicted Days Worked	17.23 (0.361)	17.30 (0.153)	16.87 (0.260)	16.28 (0.566)	16.75 (0.391)	14.98 (1.541)
Days Worked BONUS=0	1.31 (0.041)	6.96 (0.101)	1.35 (0.049)	1.174 (0.072)	12.90 (0.281)	2.63 (0.362)
Out of Sample Prediction	21.47 (0.046)	19.975 (0.164)	21.48 (0.030)	21.550 (0.060)	17.77 (0.479)	19.83 (1.725)

Note: Models I and II are estimated using maximum likelihood. Models III through VI are estimated using the method of simulated moments with an optimal weighting matrix. We report the elasticity of days worked with respect to the bonus, ϵ_{Bonus} , and the semi-elasticity with respect to a bonus cutoff, $\epsilon_{\text{Bonus_cutoff}}$. The last three rows report the expected number of days worked under the original incentives, a counterfactual where BONUS=0, and the second set of financial incentives.

Figure 4A: Predicted Fit From Model V




Figure 4B: Counterfactual Fit From Model V




Table 5: Counterfactual Cost-Minimizing Policies

Expected Days Worked	Bonus Cutoff	Bonus	Expected Cost
(1)	(2)	(3)	(4)
14	0	0	500
15	21	25	521
16	22	75	664
17	21	75	672
18	20	75	755
19	20	100	921
20	20	125	1112
21	16	225	2642
22	11	275	4604

Table 6: Teacher Performance

	Sept 2003-Feb 2006		Difference Between Treatment and Control Schools			
	Treatment	Control	Diff	Until Mid-Test	Mid to Post Test	After Post Test
	(1)	(2)	(3)	(4)	(5)	(6)
Percent of Children Sitting Within Classroom	0.72	0.73	-0.01 (0.01)	0.01 (0.89)	0.04 (0.03)	-0.01 (0.02)
Percent of Teachers Interacting with Students	1239	867	2106	643	480	983
	0.55	0.57	-0.02 (0.02)	-0.02 (0.04)	0.05 (0.05)	-0.04 (0.03)
Blackboards Utilized	1239	867	2106	643	480	983
	0.92	0.93	-0.01 (0.01)	-256766.00 (0.02)	0.01 (0.02)	-0.01 (0.02)
	990	708	1698	613	472	613

Notes: (1) Teacher Performance Measures from Random Checks only includes schools that were open during the random check. (2) Standard errors are clustered by school.

Table 7: Child Attendance

	Sept 03-Feb 06			Difference Between Treatment and Control Schools		
	Treatment (1)	Control (2)	Diff (3)	Until Mid-Test (4)	Mid to Post Test (5)	After Post Test (6)
<i>A. Attendance Conditional on School Open</i>						
Attendance of Students Present at Pre-Test Exam	0.46	0.46	0.01 (0.03)	0.02 (0.03)	0.03 (0.04)	0.00 (0.03)
	23495	16280	39775			
Attendance for Children who did not leave NFE	0.62	0.58	0.04 (0.03)	0.02 (0.03)	0.04 (0.04)	0.05 (0.03)
	12956	10737	23693			
<i>B. Total Instruction Time (Presence)</i>						
Presence for Students Present at Pre-Test Exam	0.37	0.28	0.09 (0.03)	0.10 (0.03)	0.10 (0.04)	0.08 (0.03)
	29489	26695	56184			
Presence for Student who did not leave NFE	0.50	0.36	0.13 (0.03)	0.10 (0.04)	0.13 (0.05)	0.15 (0.04)
	16274	17247	33521			
<i>C. Presence, by Student Learning Level at Program Start (for those who did not leave)</i>						
Took Oral Pre-Test	0.50	0.36	0.14 (0.03)	0.11 (0.03)	0.14 (0.05)	0.15 (0.04)
	14778	14335	29113			
Took Written Pre-Test	0.48	0.39	0.10 (0.06)	0.07 (0.07)	0.07 (0.06)	0.11 (0.07)
	1496	2912	4408			

Notes: (1) Standard errors are clustered at the level of the school. (2) Child attendance data were collected during random checks. (3) The attendance at the pre-test exam determined the child enrollment at the start of the program.

Table 8: Descriptive Statistics for Mid Test and Post Test

	Mid Test			Post Test		
	Treatment	Control	Difference	Treatment	Control	Difference
<i>A. Attrition Process</i>						
Percent Attrition	0.11	0.22	-0.10 (0.05)	0.24	0.21	0.03 (0.04)
Difference in Percent Written of Pre-Test attriters-stayers	0.01	0.03	0.02 (0.06)	0.06	-0.03	0.10 (0.06)
Difference in Verbal Test of Pre-Test attriters-stayers	0.05	0.08	-0.03 (0.14)	0.02	0.12	-0.10 (0.14)
Difference in Written Test of Pre-Test attriters-stayers	-0.41	-0.23	-0.18 (0.34)	-0.19	-0.13	-0.06 (0.29)
<i>B. Exam Score Means</i>						
Took Written	0.36	0.33	0.03 (0.04)	0.61	0.57	0.04 (0.05)
Math	0.14	0.00	0.14 (0.10)	-0.08	-0.24	0.16 (0.15)
Language	0.14	0.00	0.14 (0.10)	1.71	1.60	0.11 (0.11)
Total	0.14	0.00	0.14 (0.10)	0.35	0.24	0.12 (0.11)

Notes: (1) Test Scores in Panel B are normalized by the mean of the mid-test control. (2) Standard Errors are clustered by school.

Table 9: Estimation of Treatment Effects for the Mid- and Post-Test

Mid-Test				Post-Test			
Took Written (1)	Math (2)	Lang (3)	Total (4)	Took Written (5)	Math (6)	Lang (7)	Total (8)
<i>A. All Children</i>							
0.04 (0.03) 1893	0.15 (0.07) 1893	0.16 (0.06) 1893	0.17 (0.06) 1893	0.06 (0.04) 1760	0.21 (0.12) 1760	0.16 (0.08) 1760	0.17 (0.09) 1760
<i>B. With Controls</i>							
0.02 (0.03) 1893	0.13 (0.07) 1893	0.13 (0.05) 1893	0.14 (0.06) 1893	0.05 (0.04) 1760	0.17 (0.10) 1760	0.13 (0.07) 1760	0.15 (0.07) 1760
<i>C. Took Pre-Test Oral</i>							
	0.14 (0.08) 1550	0.13 (0.06) 1550	0.15 (0.07) 1550		0.2 (0.14) 1454	0.13 (0.09) 1454	0.16 (0.10) 1454
<i>D. Took Pre-Test Written</i>							
	0.19 (0.12) 343	0.28 (0.11) 343	0.25 (0.11) 343		0.28 (0.18) 306	0.28 (0.11) 306	0.25 (0.12) 306
<i>E. Girls</i>							
0.07 (0.03) 891	0.18 (0.07) 891	0.18 (0.07) 891	0.2 (0.07) 891	0.07 (0.05) 821	0.22 (0.12) 821	0.17 (0.09) 821	0.18 (0.09) 821
<i>F. Boys</i>							
0.02 (0.04) 988	0.12 (0.09) 988	0.14 (0.07) 988	0.14 (0.07) 988	0.05 (0.04) 929	0.19 (0.15) 929	0.16 (0.10) 929	0.16 (0.10) 929

Notes: (1) The table presents the coefficient estimate of being in a treated school on the sum of a child's score on the oral and written exams. All regressions include controls for the child's learning levels prior to the program. (2) The mid and post test scores are normalized by mid test control group. (3) Controls in Row B include Block, Teacher Test Scores, and Infrastructure Index. (4) Standard errors are clustered by school.

Table 10: Dropouts and Movement into Government Schools

	Treatment (1)	Control (2)	Diff (3)
Child Left NFE	0.44	0.36	0.08 (0.04)
Child Enrolled in Government School	0.26	0.16	0.10 (0.03)
Child Dropped Out of School	0.18	0.20	-0.02 (0.03)
N	1136	1061	2197

Notes: (1) Standard errors are clustered at the level of the school. (2) Dropouts are defined as those who were absent for the last five random checks in which a school was found open.

Table 11: Does the Random Check Predict Test Scores?

Method:	OLS	OLS	OLS	2SLS
Sample:	Control Schools	Treatment Schools	Treatment Schools	All Schools
Data:	Random Check	Random Check	Photographs	Random Check
	(1)	(2)	(3)	(4)
<i>A. Mid-test (Sept 03-April 04)</i>				
Took Written	0.02 (0.10)	0.28 (0.08)	0.36 (0.11)	0.26 (0.19)
Total Score	0.20 (0.19)	0.39 (0.21)	0.87 (0.22)	1.07 (0.43)
N	878	1015	1015	1893
<i>B. Post-test (Sept 03 -Oct 04)</i>				
Took Written	0.24 (0.16)	0.51 (0.15)	0.59 (0.20)	0.33 (0.22)
Total Score	0.58 (0.35)	1.17 (0.36)	0.98 (0.53)	0.97 (0.47)
N	883	877	877	1760

Notes: (1) The table presents the coefficient estimate of the teacher's attendance on the sum of a child's score on the oral and written exams. All regressions include controls for the child's learning levels prior to the program. (2) The mid and post test scores are normalized by the mid test control group. (3) Standard errors are clustered by school.