Teacher Incentives in the Developing World

Paul Glewwe
University of Minnesota

Alaka Holla
Innovations for Poverty Action

Michael Kremer
Harvard University

September 11, 2008

We thank Rachel Glennerster and Matthew Springer for many useful comments on this paper. This paper draws on "Increasing the Quantity and Quality of Human Capital Investment in the Developing World".

1. Introduction

In many developing counties even children who attend school seem to learn remarkably little, as evidenced by their low scores on internationally comparable tests.[1] One, though certainly not the only, explanation for this may be that teacher incentives are extremely weak in many developing countries, with teachers often having high rates of absence, despite salaries that are often far above market levels. This paper reviews education policies that attempt to improve the quality of schooling in developing countries by changing teacher incentives, including policies that improve working conditions so that teachers are more motivated to come to work; policies that provide direct payments to teachers based on either their inputs (e.g. attendance) or their outputs (e.g. student performance on tests); and policies that change teacher incentives by modifying the management of the education system, such as local hiring of contract teachers and providing information on school and student performance to the local community. This paper focuses on student achievement and teacher effort in primary and secondary education, and on evidence from randomized evaluations since they allow one to isolate the causal impact of specific policies or programs from student or school characteristics that may be correlated with those policies.

In many societies enforcement of rules requiring teachers to attend school have broken down. Providing payments to teachers based on attendance does induce more teacher attendance and does not lead teachers to show up but then not teach. However, there may be a political challenge in implementing these. Evidence from developing countries on systems that link teacher pay to the student test scores is mixed ; one study finds evidence that such policies lead to multi-tasking (teaching to the test) problems, but a recent study that is still in progress finds more positive impacts on student achievement. Turning to policies that modify the management of the education system, the limited evidence to date suggests that simply providing information to communities on teacher performance without actually changing authority over teachers has little impact. One

promising and politically feasible policy is hiring additional contract teachers locally outside the civil service system. These teachers are much cheaper than civil service teachers, and there is evidence that they attend school more often, and their students seem to do better.

The next section reviews evidence that student performance is disturbingly low in much of the developing world, and argues that teachers in many developing countries have very weak incentives. Section 3 reviews evidence that shows increases in teacher effort in response to better working conditions, while Section 4 presents results from interventions that offer more direct incentives by differentiating teachers' pay based on their performance. Section 5 discusses more fundamental reforms of education systems, based on local empowerment and contract hiring, and Section 6 concludes.

2. Background: Student Achievement and Teacher Incentives in Developing Countries

This section provides basic information on education outcomes and teacher incentives in developing countries. It begins by presenting evidence that student performance in many developing countries is much lower than in developed countries, and then it presents evidence that there are few pecuniary incentives for teachers to exert more than very minimal effort in developing countries.

A. Low Student Achievement. Children currently attending school in developing countries often receive a poor quality education, leaving school with very low levels of achievement. In Bangladesh, for example, Greaney, Khandker, and Alam (1999) found that 58 percent of rural children age 11 and older could not identify seven of eight presented letters, and 59 percent correctly answered only five or fewer of eight tasks requiring recognition of one- and two-digit numbers, writing one-digit numbers, and recognizing basic geometric shapes.[2] In Ghana, the mean score of grade 6 students on a very simple multiple-choice reading test was 25 percent, the score one would

expect from random guessing.[3] In India, 36 percent of grade 6 students were unable to understand and correctly answer the following question: "The dog is black with a white spot on his back and one white leg. The color of the dog is mostly: (a) black, (b) brown, or (c) grey".[4]

Internationally comparable achievement data are rarely available for low-income countries, but the Trends in International Mathematics and Science Study (TIMSS) does include a few, and scores in these countries are particularly low. The average eighth-grade test-taker in South Africa, for example, answered only 18 percent of questions correctly on the math portion of the TIMSS in 2003, while test-takers in the United States answered an average of 51 percent of questions correctly. The academic achievement gaps among all children of a given age are likely to be even larger than gaps among children who are currently in school, since developing countries have lower enrollment rates and low rates of progression to secondary school, so that the students in these countries who remain in school long enough to be tested in these evaluations are likely to be better than average students.

Hanushek and Wößmann (2008) provide evidence from the existing internationally comparable tests which shows that, in many developing countries, students in a given grade or of a given age score much lower than comparable students in developed countries. For example, they show that the share of students who fail to achieve basic competence is math and science (defined as scores falling between one standard deviation below the OECD mean and one standard deviation above) is quite small in almost all developed countries (usually less than 10 percent), while in many developing countries more than half of the students do not achieve this basic level of competence (examples include South Africa, Brazil, Morocco, Peru and Saudi Arabia).[5]

B. Weak Teacher Incentives. One possible cause of this lower student performance in developing countries is weak incentives for teachers to work effectively as evidenced by high rates of teacher absenteeism. In surveys of provider absence across countries and across states in India,

Chaudhury et al (2006) find that primary school teachers were absent from school 27% of the time in Uganda, 25% in India, 14% in Ecuador and 11% in Peru. They find that absence rates are typically higher in poorer countries, and that within India poorer states have higher absence rates.[6] Data from Kenya reveal a teacher absenteeism rate of 20% in rural primary schools; for comparison, absence rates among staff at a non-profit organization working in the same area are around 6%.

These rates of teacher absence seem fairly broadly distributed among the population of teachers rather than concentrated among a subset of teachers with very high absence rates. In India, for example, where the average absence rate among primary school teachers is 25 percent, if the same teachers were always absent 75 percent of teachers should never be absent in the three surprise visits made to schools, while 25% should be absent for all three visits. Instead, only 49 percent of teachers were absent zero times, while 33 percent were absent once, 14 percent twice, and only 5 percent were absent three times.[7] Chaudhury et al (2006) find a similar pattern of widespread absence in Bangladesh, Indonesia, Peru, and Uganda, but not in Ecuador. With data from two Kenyan districts, Glewwe, Ilias, and Kremer (2008) assume that the underlying probability of absence follows a beta distribution and estimate its distribution in two districts of Kenya using a maximum likelihood approach. They do this because a finite number of visits to each school causes the dispersion of absence rates in the true empirical distribution to exaggerate the underlying dispersion of probabilities of attendance among teachers. This calibration exercise shows that absences are not primarily due to a few teachers who have very high rates of absenteeism; instead, the majority of absences appear to be accounted for by teachers who attend between 50 percent and 80 percent of the time, with the median teacher absent 14 to 19 percent of the time.[8]

Even when teachers are in school, they are not necessarily teaching their students. In India, while 75 percent of teachers were present in the school, only about half of all teachers were actually teaching in the classroom when enumerators arrived.[9] Enumerators who visited classes in Kenya to

observe pedagogy found that teachers were absent from class about 27% of the time, although they were absent from school only 20% of the time. Moreover, only a small percentage of Kenyan teachers were in the classroom at the time that the class officially started. Casual observation suggests that teachers are often drinking tea in the staff room with other teachers during class time.

Public school teachers work under civil service protection in many developing countries.[10] For these teachers, effective sanctions are rarely imposed for behavior that would invite disciplinary action in developed countries. According to a survey of 3,000 headmasters of government schools in India, only one teacher in those 3,000 schools was reported to have been fired for repeated absence.[11] Incentives seem to be somewhat stronger under other contractual or organizational forms, although absence rates are high under these as well.

The same survey also looked at absence among private school teachers in India and found that, in general, their absence rates were only 2 percentage points lower than those of public school teachers. In villages with both public and private schools, however, the teacher absence rates in private schools were about 8 percentage points lower, indicating private schools are disproportionately located in villages with particularly high rates of teacher absence in public schools. One way to interpret this is that private schools tend to enter areas where public school performance is poor. Another interpretation, however, is that the entry of private schools triggers the exit of influential families from public schools and thus weakens pressure on public school teachers to attend.

Private school teachers in India are paid only about one fourth of what public school teachers are paid, and this higher attendance for lower pay may well reflect enforcement of sanctions for absence in private schools. Chaudhury et al (2006) find that 35 out of 600 private schools in their survey reported a case of the head teacher dismissing a teacher for repeated absence or

tardiness. Thus teacher dismissal for repeated absences was found to occur in about 6% of private schools, compared to only 0.03% of public schools.[12]

There is little evidence that high absence rates are due to low teacher salaries – in many developing countries teachers are paid up to five times per capita GDP. Teacher salaries across Indian states are relatively flat, so salaries for teachers in poor states are considerably higher relative to the cost of living and outside opportunities than are salaries in rich states. Absence rates are higher, however, in poor states. In many developing countries, salaries are highly correlated with the teacher's age, experience, educational background, and rank (such as head teacher status), but again there is little evidence that any of these factors is significantly associated with lower absence. In fact, across countries head teachers are significantly more likely to be absent, and point estimates suggest that better-educated and older teachers are, on average, more likely to be absent.[13] This could be due to better outside opportunities for these teachers.

Working conditions, however, do appear to motivate teachers to come to school. Absence is negatively correlated with an index that measures school infrastructure (toilets, covered classrooms, non-dirt floors, electricity, and a school library), with a one standard deviation improvement in school infrastructure associated with a 2.7 percentage point reduction in teacher absence on average across countries.[14] If frequently absent teachers are sanctioned by being transferred to schools with poor facilities, then the direction of causality from this relationship would be unclear. However, when Chaudhury et al (2006) restrict the analysis to Indian teachers on their first posting, which would have been determined by an algorithm unrelated to their performance, the relationship between absence and school infrastructure remains. Logically, if teachers' employment is not conditional on absence, pay is likely to have little effect on absence, but the attractiveness of the work environment might affect whether a teacher decides to come to work on a particular day.[15]

Overall, in developing countries students learn much less than students learn in developed countries. There is evidence from many developing countries that teacher absence rates are high and, more generally, that teachers face weak incentives to put effort into their teaching. The rest of this paper reviews evidence from randomized evaluations that assess the impact of various kinds of teacher incentive policies on student education outcomes in developing countries.

3. Incentives from the School Environment

The discussion in Section 2 suggested that the school environment can influence teachers' propensity to come to work, since teachers are more likely to be present if their schools have good infrastructure. This section examines the impact on teacher attendance of changes in the school environment that make schools more attractive places for teachers to work but that do not attempt to reward some teachers more than others based on teacher behavior or on student outcomes. The evidence, which comes from two recent randomized interventions in Kenya, suggests that teacher effort is greater when students are better prepared and motivated to learn, although, as discussed below, it is not possible to rule out other hypotheses.

After primary school fees were abolished in Kenya in 2003, the associated influx of students increased both pupil-teacher ratios and the degree of heterogeneity in incoming students' academic preparation. In response, International Child Support (ICS) Fund, a non-governmental organization (NGO), launched the Extra Teacher Program (ETP), which provided the school committees of randomly selected schools funds to hire an extra teacher for grade 1 on a contractual basis, allowing the school to add another grade 1 class. (Schools not randomly selected for the ETP serve as a comparison group.) The grade 1 children were split into two classes in two different ways. In half of the schools that were randomly selected to receive an extra teacher, students were randomly assigned to either the contract teacher or the regularly civil service teacher. In the other half, children were

placed in classrooms based on their pre-intervention achievement scores, so that children above the median level would be taught in one classroom and those below the median level in another. These classrooms were then randomly assigned to either the contract teacher or the civil service teacher.

Civil service teachers in the "tracked" schools with the extra contract teachers were around 5.4 percentage points more likely to be found in schools on a random day relative to their counterparts in the "untracked" schools with extra contract teachers. This effect seems entirely due to greater presence of teachers assigned to the "above the median" class, which may reflect easier or more pleasant teaching conditions when students are better prepared and less heterogeneous, since teachers would not have to keep track of a wide variety of different levels of learning abilities in the classroom and adjust lesson plans to accommodate a wide range of levels.

Another randomized intervention offered merit-based scholarships to the top-scoring 15 percent of sixth grade girls in Kenyan primary schools (the Girls Scholarship Program). These scholarships consisted of grants (to schools) that covered school fees, grants to parents for school supplies, and recognition in an awards ceremony. After the introduction of the program, average student test scores were 0.12 standard deviations higher in treatment schools[16], with spillover benefits for boys and for girls in the bottom of the achievement distribution (whose chances of winning the scholarships were small). Teacher attendance in program schools increased by 4.8 percentage points from a baseline attendance rate of 84 percent, which implies a 30 percent decrease in absence. Teacher behavior in the classroom, however, did not change.

Teachers may have been more motivated to show up for school because their students were more motivated. However, other explanations are also possible. Kremer, Miguel, and Thornton (2008) report anecdotal evidence that parental monitoring played a role; parents in treatment schools would visit schools more frequently to monitor teachers and pressure them to work hard so that their daughters could win scholarships.[17]

4. Rewards for Performance

A potentially more powerful, but also more politically controversial, approach to provide incentives for teachers is to differentiate teacher compensation based on teachers' inputs or outputs. Input-based incentives reward improvements in what are thought to be behaviors that improve teaching—frequent attendance or improvements in pedagogy, for example. Outcome-based incentives tie rewards to output, most frequently student learning as measured by test scores. This section examines recent evidence on both types of such incentive schemes, focusing on recent experimental studies conducted in developing countries.

A. Input-Based Incentives. In many developing countries, headmasters are supposed to record teacher attendance, and school inspectors are supposed to visit periodically and (among other things) verify that teachers are present and that teacher attendance records are accurate. This is the standard mechanism for monitoring teacher attendance. While in theory this should work, in practice it often fails because it is not enforced. Three recent studies from India and Kenya suggest that input-based rewards can be effective if they are enforced, but it may be difficult for local monitors to enforce them.

Kremer and Chen (2001) evaluate an inputs-based incentive intervention in Kenyan preschools in which teachers were eligible for bonuses equivalent to as much as 85 percent of their salary, based on their attendance, and find that the program affected neither teacher attendance nor most measures of teacher pedagogy. Moreover, it had no effect on either pupil test scores or pupil attendance. This input-based incentive program assigned headmasters to monitor teacher absence and award bonuses based on a pre-specified rule, in which bonuses were reduced for each day a teacher was absent without excuse. Any funds that were not paid to teachers went to the general school account, so headmasters had at least some financial incentive not to pay teachers bonuses if

they were not present. Nonetheless, headmasters routinely paid the entire bonus to teachers even though independent monitoring showed that teacher absences did not significantly decline relative to the baseline rate of 29%.[18]

Duflo, Hanna, and Ryan (2007) report on a randomized evaluation of an inputs-based program in remote, non-formal education centers (NFEs) in India. These non-formal schools in rural areas of India typically teach basic numeracy and literacy skills to children who are not attending formal schools, with the aim of eventually mainstreaming these children into the regular school system. The teachers at these schools do not have civil service protection, are typically paid very low wages, and have less education than regular teachers. The fact that these teachers are not part of the civil service system and were hired by an NGO, rather than the government, made it easier to implement an incentive program.

The program was implemented in a poor, sparsely populated rural area, where the NGO administering the non-formal schools found it difficult to monitor the centers. Perhaps as a consequence of this difficulty, teacher absenteeism was high (around 44 percent), prior to the intervention, despite an official policy of dismissal for absent teachers in these non-formal schools.[19]

The incentives program provided teachers cameras with tamper-proof time and date functions. A child had to photograph the teacher and the other children at the beginning and the end of each school day. Teachers' salaries were then a function of the number of "valid school days," in which the morning and afternoon pictures were separated by a minimum of 5 hours and at least 8 children appeared in each picture. Teachers received a base salary for working at least 20 days per month, and a bonus for each day in excess of 20; they were fined for each day they fell short of the minimum requirement of 20 days. Duflo, Hanna, and Ryan (2007) find that the program had a large effect on teacher absence over its 30 month duration—it decreased teacher absence in treatment schools by 19 percentage points, which is roughly equivalent to halving the absence rate.[20]

One key question with input-incentive based programs such as this one, which reward teachers for being at school but not for what they do there, is whether teachers find ways to increase the measured input without actually producing the desired educational output. For example, teachers might come to school, but then not bother to teach. Reassuringly, however, Duflo, Hanna, and Ryan (2007) find that the teachers in treatment schools were no less likely to be teaching on the days they were at school than were the teachers in comparison schools.

Further evidence that the camera program did not lead to an ineffective increase in inputs is that students scored higher on tests at the end of the program. After one year, the program increased language scores by 0.21 standard deviations, math scores by 0.16 standard deviations, and overall test scores by 0.17 standard deviations. In addition, the graduation rate to mainstream government schools increased by 10 percentage points (from a baseline of 16 percent). This improvement may simply reflect increased student time in school. While students' attendance rates (measured by unannounced spot checks) on the days their schools were open did not increase, the total amount of time children spent in school did increase since schools were open more days: children in treatment schools obtained, on average, 2.7 more days of schooling per month.[21]

Duflo, Hanna, and Ryan (2007) also show that the camera program was relatively cost-effective. The average salary was the same in treatment and comparison schools. All other costs of the program amounted to roughly $6 per child per year (assuming 20 children per teacher), which implies a cost $60 for an additional school year. In terms of raising test scores, the per child cost of increasing test scores by 0.10 standard deviations was only $3.58.

Since teacher unions often have substantial political influence and may oppose efforts to link payments to either inputs or outputs, it is worth asking whether such programs could be sustained politically in government schools. Some discouraging evidence is provided by examining an intervention among government health care workers in rural areas of Udaipur district, India.[22] In

these areas, sub-centers are facilities that provide the most basic health care services (e.g. first-aid and prenatal care) and should be staffed by at least one assistant nurse midwife (ANM). They are supposed to be open six hours per day, six days per week. A 2003 survey, however, showed that sub-centers were closed 56 percent of the time and that in only 12 percent of the cases was the ANM on duty somewhere around the center.[23]

The NGO that had administered the camera program collaborated with state and local health administrations to provide ANMs with incentives to show up to their assigned health centers on certain days on which they knew they would be monitored, by using time and day stamp machines and a specific schedule of fines and other punishments to determine an ANM's wages. During each month, if an ANM was absent for more than 50 percent of the time on the monitored days, then her pay would be reduced by the recorded number of absences for that month. If she were absent more than 50 percent of the monitored days for a second straight month, she was supposed to be suspended from government service. Some absences, such as government-mandated meetings, survey work, or other health-related work were excused and not counted as absences. If a stamp machine broke, an ANM could not be monitored until it was fixed.

In the first six months, Banerjee, Duflo, and Glennerster (2008) find that the stamp machine program doubled ANM presence, as observed in random unannounced visits to the sub-centers. On the days when they were supposed to be monitored, ANMs in the treatment group were present 60 percent of the time, while the comparison group of ANMs were present less than 30 percent of the time. The ANMs in the treatment group were found in the sub-center less than 40 percent of the time on the days when they were not expecting to be monitored.

In the next 9 month period, however, ANM presence in treatment sub-centers fell dramatically to 25 percent on monitored days, while the comparison ANMs were present 35 percent of the time, so that by the end of the period, the treatment ANMs had higher absence rates than the

ANMs in the comparison group. On the non-monitored days, the presence of treatment and comparison ANMs converged by the end of the period to a rate of less than 20 percent.[24]

The administrative data, however, did not reflect this decrease in presence. The number of days marked as exempt, as well as days with machine problems, did increase over the life of the program duration. Program monitors, in fact, reported that a number of machines had been deliberately broken and that ANMs were not present to meet the monitors when they came to sub-centers to replace the malfunctioning machines, thus extending the period in which they could not be monitored under program rules.[25]

Although one should be cautious about drawing broad generalizations from only three studies, these studies suggest that the effectiveness of input-based incentives depends on their design and implementation, in particular on the system for monitoring the inputs on which payments are based. The use of an external observation technology that produces frequent, clear-cut yes/no measurements of absence and so cannot be vulnerable to shadings of excuse (viz. cameras in India, with payments determined by distant NGO workers) appears to be more effective than using local observers, such as headmasters, who might be very close to teachers and therefore unlikely to sanction them a great deal. A simple yes/no technology with an infrequent inspection system is less likely to be effective because then a teacher who is usually absent a small fraction of the time, but who by chance is absent on the days the inspector happens to visit, would be subject to disciplinary action. Finally, the Udaipur stamp machine example suggests that, at least in some settings, it is politically difficult to sustain efforts to increase service provider attendance through more frequent monitoring, perhaps due to the political power of providers and their unions.

B. Output-Based Incentives. Linking teacher pay to student performance could potentially lead teachers to increase not only school attendance, but also to provide other inputs, such as better pedagogy, more homework, or more time spent teaching conditional on their

presence in school. In developed countries, opponents of teacher incentives based on students' test scores argue that since teachers' tasks are multidimensional and only some aspects are measured by test scores, linking compensation to test scores could cause teachers to sacrifice the promotion of curiosity and creative thinking in order to teach the skills tested on standardized exams.[26, 27]

The extremely weak teacher incentives and teacher supervision systems in many developing countries raise the potential for both these benefits and these costs if output-based incentives are implemented. On one hand, it can be argued that teachers in many developing countries are already teaching to the test and that the main problem is to get teachers to come to school, and thus teacher incentives are particularly appropriate for developing countries. On the other hand, developing countries with weak teacher accountability systems may be more prone to attempts by teachers to game any incentive system.

It is not a priori clear whether teacher incentives will improve long-run learning. Suppose teachers can exert two types of effort: genuine teaching effort to promote long-run learning and "signaling effort," which improves scores in the short-run but has little effect on long-run learning. Observable test scores depend both on underlying learning (produced by current and past teaching effort) and contemporaneous signaling effort. Test score-based incentives could potentially either increase or decrease teaching effort, since teaching effort and signaling effort could be either substitutes or complements. For example, they can be substitutes if there is a fixed amount of time in the day that must be allocated between them. On the other hand, they can be complements if there is a fixed cost to teachers of attending school at all.

Although it would be expensive and very difficult to monitor both kinds of effort on an ongoing basis, there may be ways to distinguish teaching and signaling effort empirically at the aggregate level for incentives programs that pay rewards based on test scores. First, direct observation of teacher behaviour can provide some clues. Teacher attendance or homework

assignments would most likely contribute to students' long-run learning, while other activities are likely to have a higher signalling component. In many countries, for example, some schools conduct extra test preparation or coaching outside of normal class time, often during school vacations. While these prep sessions may include some genuine learning, relative to normal classes they contain a higher proportion of effort aimed at raising test scores, such as reviewing old exams or teaching students not to leave blanks on multiple choice exams.

Second, improved learning should have a longer lasting effect on test scores than would signaling effort. Thus, a finding that test score gains do not persist beyond the end of an incentive program suggests that the program led to an increase in signaling effort, but not in effort aimed at promoting long-run learning.

A third way to distinguish efforts to increase long-run learning from test preparation activities is to examine in more detail the pattern of any test score increases. If test scores increased only on exams linked to incentives, were concentrated in subjects prone to memorization, or increased primarily on exams with formats that are more amenable to coaching, such as multiple choice exams, one can infer that test preparation activities were more common than effort to increase long-run learning.

A program in Kenya that linked teacher bonuses to student performance led teachers to modify their behavior in ways that raised scores on the tests that were used to allocate the bonuses, but did not improve teacher attendance. Although students had higher test scores on the tests tied to the incentives while the program was in effect, similar increases were not present on tests not linked to the incentives, and a year after the program ended the gains in student achievement on the tests linked to the incentives had completely dissipated.[28]

This output-based program, which was implemented in randomly selected schools in western Kenya, provided prizes to all teachers in grades 4 to 8 based on the performance of the

school as a whole on the government (district) exams that are administered in those grades each year. These prizes, which were awarded at the end of each year to top-scoring and most improved schools, ranged in value from 21 to 43 percent of the typical teacher's monthly salary. The program created incentives not only to raise test scores but also to reduce dropout rates, since all students enrolled at the beginning of the program were included in the computation of scores and, for purposes of awarding prizes, students who did not take the government exams were assigned a score less than what they could have obtained by random guessing. Calculating the school's score using only initially enrolled students also served to discourage schools from recruiting strong students to take the exams.

More students in treatment schools took the government exam than in comparison schools during the two years that the Kenyan output-based program operated (6.0 percentage points higher in the first year and 10.8 percentage points higher in the second year, from a base of 70 percent). The program did not affect dropout or repetition rates, nor did it change grade 8 graduation rates. Teachers apparently induced more students to take the exams, which improved their scores on the formula used to award teachers, but do not appear to have induced them to stay in school.

The program did generate gains in test scores, but these gains were short-lived, which is consistent with the idea that teachers were focused on activities to raise test scores in the short run rather than raise long-run learning. Students in the treatment schools scored an average of 0.14 standard deviations higher on the government exams in the second year of the program, relative to students in comparison schools, when the sample is restricted to students who took the exam in both the pre-intervention and intervention years. Yet in post-intervention years there were no significant achievement differences between the treatment and comparison schools on the government exams.

Although it is possible that achievement gains did not persist due to a natural process of depreciation (or forgetting) in what has been learned over time (Andrabi, Das, Khwaja, and Zajonc, 2008), several other patterns in the data suggest that teachers did not modify their behavior to improve long-term learning, but rather focused on short-run signaling. First, teacher absence rates were not reduced; teachers continued to be absent from school 20 percent of school days. Second, students did not report an increase in homework assignments nor did observers notice any change in pedagogy (such as increases in the use of blackboards or teaching aids, or in teachers' levels of energy or caring). Instead, schools increased their number of exam preparation sessions, especially during vacation periods. In the second year of the intervention, for example, treatment schools were 7.4 percentage points more like to conduct exam preparation sessions than were comparison schools.

Third, the achievement gains during the intervention appear to be the strongest in the subject—Geography, History, and Christian Religion—that arguably requires the most memorization. In the second year of the program, students in treatment schools scored 0.34 standard deviations higher on the government exams for this subject.

Fourth, the NGO administering the program also tested students with an exam that differed from the government exams in content and format, and which was designed to detect performance differences among a wider range of students. Whereas incentives were attached to performance on the government exam, no incentives were attached to performance on the NGO exam. Strikingly, in neither of the intervention years did students in the treatment schools display any statistically significant gains on this exam relative to their counterparts in the comparison schools.

Finally, there is evidence that the Kenyan output-based program improved students' test taking techniques. Standard test taking advice includes not leaving blanks on multiple choice questions and managing time so as to reach the end of exams. Examination of specific questions on

the NGO tests suggests that students in treatment schools were less likely to leave blanks, more likely to answer multiple choice questions, and less likely to leave answers blank at the end of the test. Together, these results suggest that the extra prep classes in program schools increased students' test taking skills. This may help explain the absence of a significant effect on the NGO exams (which were not the sort of test that prep classes were designed to teach to), particularly on questions with a format differing from that of the government tests.

A program in India that linked teacher pay to student test scores had effects that were in many ways similar to the effects of the Kenyan program, except that students also performed better on exams not linked to incentives, which suggests the program may have improved long-run learning. Muralidharan and Sundararaman (2008) examine a teacher incentive program in India that was part of the Andhra Pradesh Randomized Evaluation Study (APRESt). There were two variants of the program, a group-based teacher incentive and an individual teacher incentive. In the individual-incentives intervention, each teacher was paid 500 rupees (more than $10) for every percentage point improvement in his or her students' test scores beyond a minimum improvement of 5 percentage points. An improvement in scores of 10 percentage points would yield a bonus equivalent to about 30% of a monthly salary. In the group-incentives intervention, the same bonus was paid to all teachers based on the average performance of all students in the school. In both interventions, there were provisions to prevent gaming such as exclusion of weaker students from taking the exams.[29]

The APRESt teacher incentive program led to gains in scores on tests linked to incentives. In the first year, the gains of 0.12 to 0.19 standard deviations were similar in the individual- and school-based incentive groups. In the second year of the program, however, the individual incentives schools displayed achievement gains of 0.27 standard deviations, which were significantly

higher (at the 10 percent level) than the 0.16 standard deviation gains in the group-incentives schools.

As in the Kenyan output-based program, the APRESt teacher incentives program did not change teacher absence rates, nor did it change various indicators of classroom processes that were measured by direct observation. There is also evidence that teachers in treatment schools conducted more extra classes after school hours and gave more practice tests. Although, teachers in treatment schools were more likely to self report that after the program they assigned more homework and class work relative to their counterparts in comparison schools, these self reports should be taken with a grain of salt since the teachers in Kenya also reported assigning more homework but interviews with students suggested that they increased only the prep sessions and did not actually assign more homework. Indeed, the classroom observation data from India show no statistically significant difference between treatment and comparison schools in assigning homework, providing homework guidance or feedback, giving tests, encouraging participation, or providing help to groups or individuals. Teachers in program schools were 18 percentage points more likely to report conducting extra classes after school hours (from a base of 4 percent) and 11 percentage points more likely to report giving practice tests (from a base of 10 percent), Thus, these results appear quite similar to those from the Kenyan program: improvements in test scores coupled with increases in test preparation. Since the program is ongoing at the time of this writing, it is not possible to know whether these improvements in achievement will be sustained after the incentives have been removed.

A key difference with the results in Kenya, however, is that the APRESt teacher incentive program in India led to increases in scores not only on the tests linked to the incentives, but also on a test designed to have conceptual, rather than mechanical, questions, which had a format unfamiliar

to the students. Indeed, the India teacher incentives program appears to have generated statistically indistinguishable gains on conceptual and mechanical questions.

These results suggest that the APRESt teacher incentive program increased long-run learning rather than simply improving test taking techniques, although it is difficult to rule out the possibility that test taking techniques which helped on the mechanical test also helped on the conceptual test. It is worth noting that pupils in incentive schools scored 0.11 and 0.18 standard deviations higher in science and social studies, respectively, subjects for which there were no incentives. This is consistent either with positive spillover effects and/or an improvement in test-taking skills that could transfer across subjects. Definitive evidence on the impact of the program in India will have to await tests administered after the program is over, at which point teachers will no longer have any incentive to produce high scores. If the participants show higher test scores after the incentives end, that will provide strong evidence that the program led to actual increases in long-run learning.

In summary, the evidence on input-based systems suggests that if teachers can be given effective incentives to come to school more regularly, they will do so. In societies in which existing input-based incentive systems have broken down, and teacher absence is high, one approach to strengthening teacher incentives is to use technologies that allow monitoring by people who are less closely connected to the teachers, and in a way that provides frequent and unambiguous signals of whether teachers are present. However, the political obstacles to implementing such programs may be quite formidable.

Turning to output-based incentives, the experimental evidence from developing countries is mixed. Table 1 summarizes the results from Kenya and India. The study from Kenya suggests that linking teacher pay to a formula based on test scores led to unintended outcomes and produced little or no long-term gains in learning. The study in India has found some evidence of impacts on long-

term learning, but the program has not been in place long enough to see how long these impacts will last. It is not clear whether the differences reflect differences in the context or differences in the design of the two programs.

Table X-1: Test Score-Based Teacher Incentives

	Glewwe, Ilias, and Kremer (2008)	**Muralidharan and Sundararaman (2008)**
Context	Kenyan primary schools	Indian primary schools
Basis of prize	School average scores	School average scores, individual/teacher scores
Teacher attendance	No statistically significant effect	No statistically significant effect
Pedagogy (direct observation)	No statistically significant effect	No statistically significant effect
Test scores	0.14 SD increase	0.19 SD (math) and 0.12 SD (language)
Homework (direct observations)	No statistically significant effect	No statistically significant effect
Exam prep sessions	Increase	Increase
Effects on alternative tests that were not related to the incentives	No statistically significant effect	Yes
Evidence that test gains may reflect improvements in test taking techniques	Yes	Not discussed
Gains retained after end of program?	No statistically significant effect	Results not yet available

5. System Reforms

An alternative approach to improving teacher incentives is not simply to have more appealing working conditions or input- or output-based merit pay within the context of a civil-service system, but rather to change the lines of authority so that teachers are subject to community monitoring, local control over hiring, or a system in which schools must compete for students. Such policies have been advocated in response to the perceived failures of centralized school systems[30] and are increasingly being adopted, but at this point rigorous empirical evidence on their impact is still limited. This section reviews evidence from (a) interventions that provided communities with more information on schools in an attempt to empower parents; and (b) programs that transfer real control over hiring and firing decisions to parents.

A. Information Provision and Parental Involvement.

If lack of information or awareness about how schools should be functioning prevents people from demanding a certain standard of quality and holding schools accountable, one approach to improving social service delivery would be to provide communities with information about the quality of current schools and make them aware of their rights under national and state laws that were implemented to keep service providers accountable. A project in rural Uttar Pradesh (UP) in India sheds light on this hypothesis.

To understand the project some background information is useful. In 2000, the UP government set up Village Education Committees (VEC) in each village, which were to consist of the elected head of the village government, the head teacher of the village's government school, and three parents whose children are enrolled in the village's government school. The VECs' functions include monitoring the performance of the government school, reporting problems to higher authorities, requesting additional resources, deciding whether the contracts of the existing

community-based teachers should be renewed, recruiting new hires, and allocating any additional resources for school improvement that the school receives from a national education program.

Surveys undertaken prior to the interventions, however, show that the simple creation of local agencies does not ensure that people are informed about the state of social services or even are aware of the that such agencies exist.[31] For example, 92 percent of households surveyed did not know that a VEC existed, and among those that were aware of these committees, only 2 percent could name the members. Among the non-headmaster members of the VECs, 23 percent did not even know that they were members. Only 3.6 percent of VEC members mentioned that requesting government funds to hire a contract teacher from the community is one of the functions of the VEC. These surveys, combined with assessment tests given to children in their homes, also showed that parents considerably overestimate their children's reading and math abilities and are not fully aware of how poorly their schools are functioning.

The project in UP included three distinct interventions. In the first intervention, a team of NGO workers visited areas for two days and organized a village meeting to inform people about the current status of education in the village, the quality of local schools, state mandated provisions for schools (pupil-teacher ratios, infrastructure, midday meals, and scholarships), local funds available for education, and the responsibilities of VECs. The second intervention did all of this and also gave villagers a specific monitoring tool by actively encouraging and equipping communities to participate in testing to see whether children can read simple text and solve simple arithmetic problems. The third intervention contained all of the components of the first two and also added an option for villagers to deal with poor learning outcomes without engaging in local political institutions or the school system. It did so by adding a training component in which community volunteers with a grade 10 or grade 12 education were trained for four days to teach children how to read.

An average village of about 360 households sent about 100 people to the meetings, but all three interventions had little impact on the awareness or activism of either VECs or parents.[32] VECs in all three of the treatment groups, for example, were no more likely to perform any of their functions (such as filing complaints or hiring a contract teacher from the community) than the VECs in the comparison group. In none of the treatment groups were parents more involved with their children's school (through volunteering or filing a complaint, for example) than their comparison group counterparts, nor were they more likely to know about the state of education in their village or consider it a major issue. All three interventions also failed to improve teacher and student absence, which remained high at 25 percent and 50 percent, respectively.

The third intervention, however, dramatically improved community participation in teaching children to read and significantly improved reading achievement. In 55 of the 65 villages that received the third intervention, there were volunteers who started more than 400 reading camps with 7,500 children enrolled (roughly 8 percent of all children). When receipt of the third intervention is used to predict reading camp attendance (to avoid endogeneity bias), children who attended reading camps were 22.3 percentage points more likely to be able to read at least several letters, 23.2 percentage points more likely to be able to read at least a word or paragraph, and 22.4 percentage points more likely to be able to read a story, although this last result is not significant at the 10 percent level. Of course, these results apply only to the 8 percent of children attending the reading camps, who tended to be the weakest readers initially (13.1 percent of children who could not read even letters at baseline attended the camps, whereas only 3 percent of children who could already a story attended them.)

The results of the first two interventions together suggest that giving villagers information about the state of public goods in the village, without facilitating the use of that information, may not be very useful (the first intervention). Even when given a tool for monitoring, if community

members do not act upon the additional information from the monitoring and engage with the political institutions designed to address complaints or put pressure directly on schools, then these tools might not be effective means of improving the quality of publicly provided social services (the second intervention). There is an obvious public good/free rider problem in motivating political action to address these problems. The results under the third program suggest that moving beyond the civil-service system may have favorable results.

Two evaluations in Kenya shed light on programs that moved further in providing parents with influence, albeit limited, over teachers. As in India, it is useful to know the context. In Kenya, teachers in public primary schools are hired centrally through the teacher service commission and are assigned to schools. Their promotion is determined by the Ministry of Education, not by parents. However, Kenya also has a long standing tradition of local school committees, most of whose members are students' parents. These school committees have historically been concerned primarily with raising funds for school needs such as classroom repairs or textbook purchases. But they also potentially provide a vehicle for communication between parents and the Ministry. Moreover, school committees sometimes use some of the funds they have raised to hire contract teachers locally, whom they pay much lower wages, to supplement regular teachers provided by the Ministry of Education.

A recent program aimed to strengthen ties between school committees and local educational authorities through training and joint meetings and to improve teacher incentives with prizes awarded by school committee members. Preliminary results show that average teacher attendance did not change.[33] In treatment schools, committee members met more often with teachers, but teacher behavior in school was rarely discussed and teacher absence was never discussed. Local education officials did not increase the total number of visits to treatment schools. There is also little

systematic and significant evidence that pedagogy within the classroom changed or that student achievement improved.

Under the overall ETP program described in Section 3 which provided funds to school committees to hire these extra contract teachers, half of the school committees in untracked ETP schools and half in the tracked ETP schools were also randomly selected for training to help them monitor these contract teachers (soliciting inputs from parents, checking teacher attendance, etc.), and a formal review meeting was arranged for the committees to review the contract teachers' performance and decide whether their contract should be renewed. An evaluation showed that the monitoring program had no impact on attendance rates of the contract teacher or on the test scores of the contract teachers' students since these outcomes were statistically indistinguishable for contract teachers in the ETP schools without extra school committee monitoring and the contract teachers in the ETP schools with extra school committee monitoring.[34] This might have occurred because the contract teachers already had very high attendance rates, as discussed below.

However, the program may have had some impact on the attendance rates of the regular civil service teachers and on the test scores of their students. Training in monitoring increased the attendance rates of civil service teachers in ETP schools by 7.3 percentage points relative to civil service teachers in ETP schools whose school committees did not receive the monitoring training, although this result is not quite significant at the 10 percent level. .

The monitoring program also increased the attendance rates of students assigned to the civil service teacher by 2.8 percentage points relative to the pupil attendance trends experienced by comparable civil service teachers ETP in schools without the training program. Students assigned to civil service teachers in schools with trained committees also experienced test score gains of 0.18 standard deviations in math relative to students assigned to civil service teachers in ETP schools without the training program. Overall it seems that the program had a modest effect on teacher

incentives. A number of questions remain open, however, including whether the effect persists or is only a short run effect, and whether it generalizes beyond the particular context here in which all schools had contract teachers.

B. Local Hiring of Contract Teachers. A number of countries have moved towards local hiring of contract teachers to supplement civil service teachers. These contract teachers are usually paid a fraction of what civil service teachers are paid, often only a quarter as much. They are often hired locally rather than through the central government. Often they have lower formal educational qualifications than are required of regular teachers. The move towards hiring contract teachers is motivated partly by concern about the weak incentives faced by civil service teachers and partly by fiscal concerns since these teachers are much less expensive than regular teachers. In the cross-country study of absence discussed earlier, Chaudhury et al (2006) find that teachers from the local area are less likely to be absent in all six countries and that this association is statistically significant for India and Indonesia.[35] Limited experimental evidence exists on the effectiveness of these teachers. [36]

Evidence from India suggests that hiring outside of the civil service system can be an effective way to improve the quality of teaching. The Balsakhi program implemented in India paid young women from the community, who had completed only secondary school (grade 10), roughly one-tenth of a regular teacher's salary to teach basic literacy and numeracy skills to children who had reached grade 3 or 4 but had not yet mastered basic skills;[37] these children left the classroom and received tutoring for two hours per day. Banerjee, Cole, Duflo, and Linden (2007) find that the Balsakhi program increased average test scores in treatment schools by 0.14 standard deviations in the first year and by 0.28 standard deviations after two years, with most of this increase due to large gains experienced by children at the bottom of the initial test-score distribution and by the children who received the remedial instruction. One year after the program ended, the average student in the

program schools had a 0.1 standard deviation advantage over the average student in the comparison schools, suggesting that at least some of the effect of the program persisted.[38]

As discussed in Section 3, the Kenyan ETP intervention provided funds to a randomly selected set of school committees for hiring a local contract teacher. In this project, the contract teachers have the same academic qualifications as regular teachers.[39] Unlike their civil service counterparts, local contract teachers could be fired by school committees and were paid only a fraction of the civil service pay. The probability that contract teachers were in class and teaching was roughly 16 percentage points higher than that for civil service teachers in comparison schools, who attended 58.6 percent of the time, and 29.1 percentage points higher than that of civil service teachers in the ETP schools.

In the absence of the monitoring program or tracking of students by initial achievement, the presence of these contract teachers increased absence among the civil service teachers in ETP schools. The probability that a civil service teacher was in class and teaching in the ETP schools fell by 12.9 percentage points relative to civil service teachers in the comparison schools. Their classes would sometimes be combined with a class whose teacher was present, but at other times students would simply sit unsupervised in their classrooms. However, training school committees in monitoring and tracking students by initial achievement prevented a statistically significant decline in attendance by civil service teachers from the hiring of local contract teachers.

Overall the ETP program raised student test scores by 0.22 standard deviations. Students assigned to contract teachers scored 0.23 standard deviations higher and attended school 1.7 percentage points more often (from a baseline attendance rate of 86.1 percent) than students who had been randomly assigned to civil service teachers.

In another program that was conducted in India alongside the APRESt teacher incentives program discussed in Section 4, one set of schools also randomly received a contract teacher for two

years. Unlike Kenya's ETP program, the APREStcontract teacher program did not assign the contract teachers to particular classrooms and did not employ contract teachers with similar academic qualifications as their civil service counterparts. For example, 72 percent of contract teachers were female, and their average age was 27, while only 34 percent of their civil service teacher counterparts were female and their average age was 40. Only 44 percent of contract teachers had at least a college degree, and only 8 percent had received a formal teacher training degree or certificate, while 85 percent of civil service teachers had a college degree, and 99 percent had received a formal teacher training degree or certificate. These contract teachers were also more local: 85 percent lived in their assigned school's village, compared to only 11 percent of civil service teachers.

For the first year of the APREStcontract teacher program, Muralidharan and Sundaraman (2008) find that contract teachers were 11 percentage points less likely to be absent than their civil service counterparts in the same school, who were absent 24 percent of the time; in the second year, the contract teachers were 17 percentage points less likely to be absent than the civil service teachers, who were absent 29 percent of the time. Unlike the ETP program in Kenya, this contract teacher program in India did not induce greater absence of the civil service teachers. Supplying funds for contract teachers significantly decreased the rate of multi-grade teaching and the effective class size in treatment schools. In both years of the program, contract teachers were also significantly more likely than civil service teachers to be teaching during random spot checks at the schools.[40]

Contract teachers with lower absence rates and higher rates of teaching activity in the first year of the APREStcontract teacher program were more likely to have had their contracts renewed in the second year. Since good performance seems to translate into contract renewal in this context, this potential for renewal could have been an important incentive for the contract teachers.

Muralidharan and Sundaraman (2008) also find that the APRESt contract teacher program improved test scores. After two years, students in schools assigned a contract teacher scored 0.11 standard deviations higher than their counterparts in the same grade in schools without contract teachers.[41]

It is not clear which characteristics of contract teachers were responsible for the success in Kenya and India. Contract teachers were more likely to come from the local area than regular civil service teachers, which could have increased their accountability to their schools' parents since they might have faced effective sanctions outside of the school environment. The contract teachers in Kenya and India were also chosen by school committees rather than by a central bureaucracy, so they might have been better matched to the schools than civil service teachers. Finally, the observed differences between contract and civil service teachers could also have resulted from differences in the nature of their contracts – namely, that contract teachers were not guaranteed employment security and could be fired.

6. Conclusion

Many students in developing countries score very low on standardized tests, and one potential explanation is the weak incentives faced by teachers, as evidenced by their very high absence rates. In many developing countries systems for monitoring and enforcing teacher presence have broken down. Higher pay alone is therefore unlikely to induce acceptable attendance, although teacher attendance does appear to be responsive to working conditions. This paper has discussed a number of policies that attempt to increase student performance in developing countries by increasing teacher incentives, and Table 2 summarizes their results.

While more research is needed before generalizations can be drawn with confidence, several themes seem to emerge form these studies. There is evidence that if teachers do face strong, well designed incentives to reduce their high absence rates, they will be absent less often and students

will learn more. Using technologies that provide frequent and unambiguous signals of teacher presence to monitors who are more distant from the teachers may be a promising policy. However, there could be political obstacles to putting such systems in place. There is mixed evidence on the impact of more ambitious programs, which differentiate teacher pay through formulas based on test scores attained by students. More experimentation with these approaches is warranted, but such experimentation should be accompanied by careful evaluation so that programs that are successful can be scaled up; while unsuccessful programs can be modified or terminated.

More fundamental changes to the incentive system also seem promising. While simply providing information to communities about the performance of students and the absence rates of teachers may have a limited impact, empowering local communities to hire teachers on a contract basis can have a significant impact. It is worth noting, however, that many of these teachers are subsequently given jobs in the formal system, and that may be an important part of their motivation to work hard, so one should not necessarily assume that moving to a system in which there were only contract teachers and no civil service positions would have the desired incentive effects. That, however, is not on the table politically in many countries, and the relevant question is whether it would be better to have a system in which teachers first worked on a contract basis for several years before being given civil-service positions. The evidence presented here suggests that that would be a clear improvement. Aside from the potential incentive benefits documented here, there may also be selection benefits to the extent that only those who perform well as contract teachers are given civil-service positions.

Recent research in the U.S. shows that there are large differences in the test score gains associated with different teachers within schools, but indicators such as teacher certification, years of experience, and master's degrees have little predictive power for these achievement gains.[42] Yet Gordon, Kane, and Staiger (2006) present evidence that suggests that school districts can predict

which teachers are effective in generating test score gains, and which are not, with only two years of student outcome data. Given these two empirical findings, they recommend reducing the barriers to entering the teaching profession for teachers without traditional teacher certification and requiring a two year trial period once teachers are hired, after which they can receive tenure based on their performance in the classroom.[43] In order to work, however, these policies would require a reliable data infrastructure to measure teacher performance.

A policy of first hiring teachers on a contract basis, and then awarding civil-service status to those who are successful, could be politically viable because it offers the opportunity to expand the ranks of teachers in response to rising enrollment at modest cost but does not threaten the rents of existing civil-service teachers or threaten the long-run viability of teachers' unions. In contrast, programs which differentiate pay among civil-service teachers based on either attendance or student test scores are likely to run into opposition from teacher's unions.

A few final comments concern a policy direction which this paper has not examined, but which would perhaps generate the most fundamental change in teacher incentives – expansion of school choice under a system in which schools compete for students. Experimental results from Colombia's PACES program, which distributed private school vouchers to students from poor neighborhoods, demonstrate substantial gains in both student attendance and achievement.[44] [45] Yet non-experimental studies from Chile's national voucher program have yielded mixed results, with Gallego (2006) arguing that there are student achievement benefits in both public and private schools and Hsieh and Urquiola (2006) arguing that student achievement did not improve and that increased school choice led to increased sorting by parental income and education background.[46] [47]

The PACES studies, however, did not examine whether these gains resulted from changes in teacher behavior, and there is some evidence to suggest that the program induced greater effort from students since students who won a voucher through the program's lottery worked fewer labor

hours in subsequent years than did students who did not win a voucher in the lottery. Better matching of students to schools might also account for the observed achievement gains. More experiments in expanding school choice, with special efforts to collect data on teacher attendance and other indicators of effort, will be necessary to determine whether such policies can improve teacher incentives.

Table 2: Summary of impacts of teacher incentive programs, assessed by randomized trials

Policy	Country	Results	Source
Incentives from the school environment			
Tracking by initial test scores	Kenya	▪ Attendance of civil service teachers 5.4 percentage points higher in tracked classes relative to civil service teachers in untracked classes ▪ Effect concentrated in classrooms with "above the median" students	Dupas, Duflo, and Kremer (2007)[48]
Merit scholarships for students	Kenya	▪ Increased teacher attendance rates by 4.8 percentage points from baseline of 84%	Kremer, Miguel, and Thornton (2008)[49]
Input-based incentives			
Cameras use to verify teacher attendance and teachers paid for higher attendance in non-formal schools	India	▪ Teacher absence drops by 19 percentage points from baseline of 42% ▪ No effect on pedagogy in the classroom ▪ Students' test scores increase by 0.17 standard deviations ▪ Graduation rate to mainstream government schools increased by 10 percentage points from a baseline of 16%	Duflo, Hanna, and Ryan (2007)[50]
School headmasters monitor teacher attendance and teachers paid bonuses for higher attendance	Kenya	▪ No increase in teacher attendance ▪ Headmasters paid bonuses regardless of attendance	Kremer and Chen (2001)[51]
Output-based incentives			
Teachers received prizes for higher (school average) student test scores	Kenya	▪ Test scores increased by 0.14 standard deviations on tests linked to the incentives ▪ Gains highest in subject requiring the most memorization ▪ Gains not sustained one year after program	Glewwe, Ilias, and Kremer (2008)[52]

		▪ No gains on tests not linked to the incentives ▪ Improvement in students' test-taking techniques ▪ Increase in exam preparation sessions held by teachers but no other changes in pedagogy ▪ No decrease in teacher absence	
Teachers received prizes for higher (school average and teacher specific) student test scores	India	▪ Test scores increased by 0.12-0.19 standard deviations in first year and by 0.22 standard deviations in second year. ▪ Test scores also increased on non-incentivized subjects ▪ Test score gains also made on questions with unfamiliar format ▪ No change in pedagogy, as measured by direct observation ▪ Too early to tell if gains sustained because program is ongoing.	Muralidharan and Sundararaman (2008)[53]
Information provision and parental involvement			
Three interventions: *(i) Villagers informed about local school quality and responsibilities of local leaders* *(ii) Same as above + parents taught how to test children* *(iii) Same as (ii) + training of community volunteers to teach children to read at reading camps*	India	▪ No effect of pure information intervention ▪ No effect of training parents to test children at home ▪ 8% of all children attended reading camps; these children were 22.3 percentage points more likely to be able to read at least letters, 23.2 percentage points more likely to be able to read at least at least a word or paragraph, and 22.4 percentage points more likely to be able to read at least a story.	Banerjee, Duflo, and Glennerster (2008)[54]
Training parent school committees to monitor contract teacher	Kenya	▪ Increased attendance of civil service teachers by 7.3 percentage points relative to	Dupas, Duflo,

		civil service teachers in non-monitored schools ▪ Increased attendance of students assigned to the civil service teacher by 2.8 percentage points relative pupil attendance gains in non-monitored schools ▪ Increased math test scores of students of civil service teachers by 0.18 standard deviations ▪ No effect on contract teachers or their students	and Kremer (2007)[55]
Local hiring of contract teachers			
Young women from community paid to give remedial education to some students for 2 hours/day	I India	▪ Increased average test scores by 0.14 standard deviations in first year, by 0.28 standard deviations in second year ▪ Largest gains by children at bottom of initial test score distribution and by children receiving remedial education ▪ Average gain of 0.1 standard deviations persists one year after program ended	Banerjee, Cole, Duflo, and Linden (2007)[56]
Hiring of local contract teachers to split first grade classes	Kenya	▪ Probability of being in class and teaching drops by 12.9 percentage points for civil service teachers from baseline of 58.6% ▪ Contract teachers roughly 16 percentage points more likely to be in class and teaching relative to civil service teachers in comparison schools, and 29 percentage points more likely relative to civil service teachers in program schools	Dupas, Duflo, and Kremer (2007)[57]
Hiring of local contract teachers	India	▪ After one year, absence rates of contract teachers 10.6 percentage points less than absence rates of civil service teachers in schools without contract teachers,	Muralidharan and Sundaramanan (2008)[58]

		and 24.1 percentage points less than civil services teachers in the same school ▪ After second year, absence rates of contract teachers 17 percentage points less than absence rates of civil service teachers in schools without contract teachers, and 29.4 percentage points less than civil service teachers in the same school ▪ No significant increase in absence for civil service teachers ▪ Decrease rate of multi-grade teaching ▪ Decreased class size

Note: All cited results are statistically significant.

Works Cited

1. Eric A. Hanushek and Ludger Wößmann, "The Role of Education Quality in Economic Growth." Policy Research Working Paper 4122 (World Bank, 2007).

2. Vincent Greaney, Shahidur R. Khandker, and Mahmudul Alam, *Bangladesh: Assessing Basic Learning Skills* (Dhaka, Bangladesh: University Press, 1999).

3. Paul Glewwe, *The Economics of School Quality Investments in Developing Countries: An Empirical Study of Ghana* (New York: St. Martin's Press, 1999).

4. Marlaine Lockheed and Adriaan Verspoor, *Improving Primary Education in Developing Countries* (Oxford University Press, 1991).

5. Hanushek and Wößmann, "The Role of Education Quality in Economic Growth".

6. Nazmul Chaudhury and others, "Missing in Action: Teacher and Health Worker Absence in Developing Countries," *Journal of Economic Perspectives,* vol. 20 (2006), pp.91-116.

7. Ibid.

8. Paul Glewwe, Nauman Ilias, and Michael Kremer, "Teacher Incentives," mimeo (1999).

9. Chaudhury and others, "Missing in Action".

10. Civil service protection refers to job security provisions that are associated with being a government employee, which are often supported by membership in a labor union. Such protections include job security, relatively high wages, the right to strike, and pensions.

11. Ibid.

12. Ibid.
13. Ibid.

14. Ibid.

15. Ibid.

16. Michael Kremer, Edward Miguel, and Rebecca Thornton, "Incentives to Learn," *Review of Economics and Statistics* (Forthcoming).

17. Ibid.

18. Michael Kremer and Daniel Chen, "Interim Report on a Teacher Incentive Program in Kenya," mimeo (2001).

19. Esther Duflo, Rema Hanna, and Stephan Ryan, "Monitoring Works: Getting Teachers to Come to School," NBER Working Paper 11880 (Cambridge, Mass.: National Bureau of Economic Research, 2007).

20. Ibid.

21. Ibid.

22. Abhijit Banerjee, Esther Duflo, and Rachel Glennerster, "Putting a Band-Aid on a Corpse: Incentives for Nurses in the Indian Public Health Care System," mimeo (2008).

23. Ahhijit Banerjee, Angus Deaton, and Esther Duflo, "Health Care Delivery in Rural Rajasthan," *Economic and Political Weekly* (February, 2004), pp. 944-949.

24. Banerjee, Duflo, and Glennerster, "Putting a Band-Aid on a Corpse".

25. Ibid.

26. Bengt Holmstrom and Paul Milgrom, "Multi-Task Principal Agent Analysis: Incentive Contracts, Asset Ownership, and Job Design," *Journal of Law, Economics, and Organization,* vol. 7 (1991), pp. 24-52.

27. Jane Hannaway, "Higher Order Thinking, Job Design, and Incentives: An analysis and Proposal," *American Education Research Journal*, vol. 29 (1992), pp. 3-21.

28. Paul Glewwe, Nauman Ilias, and Michael Kremer, "Teacher Incentives," mimeo (2008).

29. Karthik Muralidharan and Venkatesh Sundararaman, "Teaching Incentives in Developing Countries: Experimental Evidence from India," mimeo (2008).
30

31. Abhijit Banerjee, and others, "Can Informational Campaigns Spark Local Participation and Improve Outcomes: A Study of Primary Education in Uttar Pradesh, India," World Bank Policy Working Paper 3967 (The World Bank, 2006).

32. Abhijit Banerjee, and others, "Pitfalls of Participatory Programs: Evidence from a Randomized Evaluation in Education in India," mimeo (2008).

[33] Joost de Laat, Michael Kremer, and Christel Vermeersch, "Teacher Incentives and Local Participation," mimeo (2008).

34. Esther Duflo, Pascaline Dupas, and Michael Kremer, "Peer Effects, Pupil-teacher Ratios, and Teacher Incentives," mimeo (2007).

35. Chaudhury and others, "Missing in Action".

[36] For non-experimental results, see Emiliana Vegas and Joost de Laat, "Do differences in Teacher Contracts Affect Student Performance," World Development Report 2004 Background Paper (2005); Emiliana Vegas, "Teacher Labor Markets in Developing Countries," *The Future of Children*, vol. 17 (2007), pp. 219-232; Ilana Umansky and Emiliana Vegas, "Inside Decentralization: How Three Central American School-based Management Reforms Affect Student Learning Through Teacher Incentives," *World Bank Research Observer*, vol. 22 (The World Bank, 2007), pp. 197-215; Emmanual Jimenez and Yasuyuki Sawada, "Do Community-Managed Schools Work? An Evaluation of El Salvador's EDUCO Program," *World Bank Economic Review,* vol. 13 (The World Bank, 1999), pp. 415-441.

[37] *Balsakhi* means "child's friend."

38. Abhijit Banerjee and others, "Remedying Education: Evidence from Two Randomized Experiments in India," *Quarterly Journal of Economics*, vol. 122 (2007), pp. 1235-1264.

39. Duflo, Dupas, and Kremer, "Peer Effects, Pupil-teacher Ratios, and Teacher Incentives".

40. Muralidharan and Sundararaman, "Teaching Incentives in Developing Countries: Experimental Evidence from India".

41. Ibid.

42. Eric A. Hanushek and Steven G. Rivkin, "How to Improve the Supply of High Quality Teachers," in *Brookings Papers on Education Policy: 2004*, edited by Diane Ravitch (The Brookings Institution, 2004).

43. Robert Gordon, Thomas J. Kane, and Douglas O. Staiger, "Identifying Effective Teachers Using Performance on the Job," The Hamilton Project Discussion Paper 2006-01 (The Brookings Institution, 2006).

[44] Angrist, Joshua and Victor Lavy, "Using Maimonides' Rule to Estimate the Effect of Class Size on Children's Academic Achievement," *Quarterly Journal of Economics,* 114(1999): 533-576.

[45] Angrist, Joshua, Eric Bettinger, Erik Bloom, Elizabeth King and Michael Kremer, "Vouchers for Private Schooling in Colombia: Evidence from a Randomized Natural Experiment," *American Economic Review,* 92(2002): 1535-1558.

[46] Gallego, Francisco, "Voucher-School Competition, Incentives, and Outcomes: Evidence from Chile," mimeo (2006)

[47] Hsieh, Chang-Tai and Miguel Urquiola, "The effects of Generalized School Choice on Achievement and Stratification: Evidence from Chile's School Voucher Program," *Journal of Public Economics*, 90 (2006): 1477-1503.

48. Duflo, Dupas, and Kremer, "Peer Effects, Pupil-teacher Ratios, and Teacher Incentives".

49. Kremer, Miguel, and Thornton, "Incentives to Learn".

50. Duflo, Hanna, and Ryan, "Monitoring Works".

51. Kremer and Chen, "Interim Report on a Teacher Incentive Program in Kenya".

52. Glewwe, Ilias, and Kremer, "Teacher Incentives".

53. Muralidharan and Sundararaman, "Teaching Incentives in Developing Countries: Experimental Evidence from India".

54. Banerjee, Duflo, and Glennerster, "Putting a Band-Aid on a Corpse".

55. Duflo, Dupas, and Kremer, "Peer Effects, Pupil-teacher Ratios, and Teacher Incentives".

56. Banerjee and others, "Remedying Education".

57. Duflo, Dupas, and Kremer, "Peer Effects, Pupil-teacher Ratios, and Teacher Incentives".

58. Muralidharan and Sundararaman, "Teaching Incentives in Developing Countries: Experimental Evidence from India".