

Addressing Absence

Abhijit Banerjee and Esther Duflo

Absent providers are a major problem both for public health facilities and primary schools in many developing countries. The paper by Chaudhury and others in this issue provides new and systematic evidence on the rates of absenteeism based on surveys of absence rates of teachers and health workers in several developing countries. For example, in India, absence rates for teachers are over 24 percent, and for health providers they are over 40 percent.

Other recent surveys confirm these findings. For example, Banerjee, Deaton and Duflo (2004a, b) analyze the results of a detailed survey of 143 government health facilities in rural Udaipur district, in the state of Rajasthan in India. In this survey, absence was monitored weekly for a year (at unannounced visits on random days), by a locally hired person, and also monthly, by a monitor who was part the research team. The monitor looked for the nurse either in her center or in any of the villages where she was supposed to be working. The average absence rate in the primary health care facilities (the larger centers) was 36 percent. The average absence rate is even higher at 45 percent in rural subcenters, and since these subcenters are generally staffed by only one person, absence means that no one is running the subcenter. Another survey in Udaipur focused on 60 nonformal education centers run by a nongovernment organization. In this case, absence was monitored monthly by a monitor who was part of the research team. The average absence rates for teachers in the these education centers was 36 percent (Duflo and Hanna, 2005). The statistics reported here concern the comparison group of this experiment.

Moreover, these year-long surveys reveal that absence is not concentrated among a few “bad apples.” While some facilities are clearly worse than others, it is

■ *Abhijit Banerjee and Esther Duflo are both Professors of Economics and Directors of the Poverty Action Lab, Massachusetts Institute of Technology, Cambridge, Massachusetts.*

not the case that a few facilities account for all the absences. For example, absences in the worst quartile of the health subcenters account for only 36 percent of the overall absences. The best quartile still accounts for 14 percent of the absences. In addition, absence patterns are erratic: there is no particular time, day or month, when one can be more or less sure that the center will be open. It is unlikely, moreover, that the villagers have much information that we do not have: the nurse has no way of informing the villagers when she is coming, and there is no evidence that she tries. As a result, villagers cannot plan around the absences. Someone who is planning to go to the health center has to decide to spend the half an hour or so that it takes to walk the 1.4 miles that separates the average village in our sample from the closest public health facility, knowing that there is only a one in two chance of seeing someone. It is no wonder then that people rarely use the public health facilities. Data from Banerjee, Deaton and Duflo (2004a) suggest that only about a quarter of the visits to health facilities in rural Udaipur district are to government facilities. The rest are to a combination of private providers, who are often entirely unqualified (fewer than 40 percent of them have a medical degree, and almost 20 percent have not completed secondary school), and traditional healers.

Efforts to improve attendance are therefore crucial to making public services play their designated role in the lives of the poor. Initiatives to reduce absence rates in schools range from hiring more teachers on short contracts and instituting school committees to decentralizing of education to local government. Unfortunately, it is rarely clear whether initiatives to fight absence are having their desired effect. Imagine, for example, that a new headmaster arrives at a school, full of enthusiasm and new ideas and wants to fight teacher absence. He gets parents involved by setting up a parents' committee. Word spreads that the new headmaster is good, and some children transfer into the school from other local schools. Teachers start showing up more regularly. How can one disentangle the effects of the parents' committee, the impact of the headmaster's enthusiasm on other teachers in school and the effect of the influx of new students, who might be more motivated than average?

The cleanest and clearest way to get around this problem and to identify the effect of a reform on attendance is to run a randomized trial. Choose 100 representative schools, establish (for example) a parents' committee in half, and wait to start such committees in the other half till the evaluation is over: over the evaluation period, compare the outcomes in the two groups of schools. This randomized evaluation approach is very similar to how new drugs and vaccines are tested, and it is now being increasingly used by development economists.

This paper discusses evidence on a number of innovative strategies to reduce absenteeism in government and nongovernment organization-run schools and health facilities that have been implemented in Kenya and India over the past few years and that have all been evaluated using the randomized evaluation methodology. These strategies have involved alternative levers to fight absence. Some have tried to improve incentives for providers, either through rewards and punishments

implemented by external monitors, or through facilitating a more active involvement of those who expect to benefit from the service. Others base their strategies on the idea that the providers are discouraged by the lack of interest among the potential beneficiaries in what they are being offered and have tried to increase the demand for the services as a way of putting more pressure on the providers. The results of these efforts, taken together, shed light not only on ways to address the problem of absence in the public sector, but also on the underlying reasons for this phenomenon.

External Control

The obvious method to fight absence is to monitor more intensively and to base incentives (both rewards and punishments) on measured performance. To do this, the employer has to establish a set of more or less explicit rules and put people in charge of the monitoring. The control is external when it is exercised by those who have no direct stake in the service being delivered. This is the case, for example, when the state seeks to punish shirking teachers or when it gives them rewards based on school attendance and/or school performance. External control need not always be about monetary incentives—the state can also use praise or shame to put pressure on its agents.


The most common type of external control is one where someone in the institutional hierarchy (like the headmaster of a school) is given the task of keeping an eye on the teacher and penalizing absences. An alternative may be to use some impersonal method, such as a camera to record absence, and then to base rewards or penalties on that data. A person responsible for monitoring may either be too lazy to monitor or might collude with workers. On the other hand, impersonal monitoring makes no allowances for the circumstances of the absence. Of course, in both cases, someone needs to be in a position to enforce the rewards and penalties linked to absence. A final alternative is that rather than measuring absence, some other measure of performance, such as test scores, which reflects more directly the final outcome we care about, could be used.

Impersonal Monitoring by Camera

A randomized experiment using impersonal monitoring was implemented by Seva Mandir, a nongovernment organization that runs nonformal single-teacher primary education centers in tribal villages in the rural Udaipur district. The program was evaluated by Duflo and Hanna (2005). Udaipur is a sparsely populated, arid and hilly region, where villages are remote and access is difficult, which makes it very difficult for Seva Mandir to monitor the education centers regularly. Absence rates are very high, despite the organization's policy calling for dismissal in cases where absence is high. At the baseline of this study in August 2003, the absence rate was 44 percent, which was quite similar to the 40 percent absence rate found in a study conducted in 1995 (Banerjee, Deaton and Duflo, 2004b). Most

Figure 1

Monitoring By Camera: A Teacher and His Students




Seva Mandir schools are one-teacher schools, such that when the teacher is absent, children just go back home and lose the entire day of schooling.

Seva Mandir selected 120 schools to participate in the study. In 60 randomly selected schools (the “treatment schools”), the organization gave the teacher a camera, with instructions to take a picture of himself or herself and the students every day at opening time and at closing time. The cameras had a tamper-proof date and time function. Figure 1 shows a sample picture of a teacher and his students; the day of the month and the time of day are indicated on the right corner of the picture. Similar pictures were received twice a day for each school that was open on that day. Teachers received a bonus as a function of the number of “valid” days they actually attended. A “valid” day was defined by a day where the opening and closing pictures were separated by at least five hours and a minimum number of children were present in both pictures. A teacher received a salary of 1,000 rupees (Rs) monthly if he was present at least 21 days in a month. Each additional valid day carried a bonus of Rs 50 (\$1 U.S. at the exchange rate, or \$6 in purchasing power parity), up to a maximum of Rs 1,300 per month. Each day missed relative to the 21-days benchmark carried a penalty of Rs 50. Therefore, the way the bonus was set up, a teacher’s monthly salary could range from Rs 500 to Rs 1,300 per month. In the remaining 60 schools (the “comparison schools”), teachers were paid Rs 1,000 and were told (as usual) that they could be dismissed for poor performance. One unannounced visit every month was also made to measure teacher absences in the comparison schools.

Figure 2

Impact of the Cameras

Number of Schools Found Open Times in Treatment and Comparison Schools
(out of 13 visits)



The program resulted in an immediate improvement in teacher attendance, which persisted during the entire year. The absence rate of teachers was cut by half in the treatment schools, dropping from an average of 36 percent in the comparison schools to 18 percent in the treatment schools. Figure 2 shows the observed density of absence rates in treatment and in comparison schools for the 13 visits that took place during the duration of the program. The program changed the shape of the distribution of absence quite radically. In comparison schools, just one teacher is present on all 13 days when he was observed. In treatment schools, 11 (out of 60) are present all 13 times, 27 are present 12, 11 and 10 times, and everyone is present at least 7 times. The camera program was effective on two margins: it eliminated extremely delinquent behavior (less than 50 percent presence), and it increased the number of teachers with a perfect or very high attendance record.

When in school, teachers were as likely to be teaching in treatment as in comparison schools, and the number of students present was the same. However, because there were fewer teacher absences, treatment schools taught the equivalent of 88 children-days more per month than a comparison school—a one-third increase in the number of child-days.¹

Given the structure of payments, the average salary in the treatment schools ended up matching almost exactly the average salary in the comparison schools. The incentives were therefore effective without an increase in the teacher's net pay: the only cost of running the program was the cost of the cameras and the administration of the program. Moreover, this cost is quite reasonable, compared to the salary of a teacher: our cost-benefit analysis reveals that the program cost only

¹ A child-day is defined as the number of children present in a given day where the school is open, or zero if the school is closed (when the teacher is absent, the school is closed, and children go back home).

\$6 per child per year to be administered (for an increase of over 30 percent in the number of days the child is taught). Most of these costs will be lower when digital cameras (which do not require developing film) can be used. The program is therefore not too expansive to be, in principle, scaled up to an entire school system.

This program shows that a straightforward incentive program, mechanically implemented, is a very effective way to reduce absence in schools. This combination of a clearly defined task, simple incentives and a systematic implementation can be quite powerful. Acceptability of the cameras among Seva Mandir teachers was high: many commented that they felt it was now in their own hands to increase their salary. Today, several months after the end of the study, the program is still in place at Seva Mandir. However, in practice, in most school systems, incentive systems are not implemented in this way, but are mediated by people in the school hierarchy, either inspectors or school headmasters.

How Mediated Incentives Lose Bite

A program implemented by a nongovernment organization called ICS Africa in Kenya suggests that when headmasters implement incentives, the incentives might lose their power. ICS Africa introduced an incentive program for preprimary school teachers.² The headmaster was entrusted with monitoring the presence of the preprimary school teacher. At the end of the term, a prize (a bicycle) was offered to teachers with a good attendance record. If a teacher did not have a good attendance record, the money would remain with the school and could be used on whatever the headmaster and the school committee preferred.

Kremer and Chen (2001) report on the results of this experiment. In all treatment schools, the headmasters marked the preschool teachers present a sufficient number of times for them to receive the prize (and they therefore all got it). However, when the research team independently verified absence through unannounced visits in both treatment and comparison schools, they found that the absence rate was actually at exactly the same high level in treatment and in comparison schools. Either to avoid the unpleasantness of a personal confrontation, or out of compassion for the preschool teachers, headmasters apparently cheated to make sure that the preschool teachers could get the prizes.

This outcome suggests that when human judgment is involved in a system where rules are often bent, incentives may easily be perverted, either as in this case in an equitable direction or to favor some specific individuals or groups. Mechanically implemented systems, such as the camera program, are immune to these problems—which may be the source of their effectiveness. Of course, they

² ICS Africa is the African branch of a Dutch organization that has implemented programs to improve school quality in western Kenya for over 10 years and works in close collaboration with the school system. For 10 years, ICS Africa has also conducted many randomized evaluations in collaboration with Michael Kremer and his co-authors. The evaluation costs were typically financed with research grants, while ICS provided funding for the inputs. This paper reports on many of their findings.

require the willingness of the administration to continue implementing them, possibly in the face of resistance by the teachers or health providers, which may eventually require either enough demand for the service or political will.

Rewards for Performance Rather than Presence

Glewwe, Ilias and Kremer (2003) describe the results of an attempt in Kenya to provide incentives to teachers based on the test scores of students in their class. ICS Africa provided prizes to teachers in grades 4 to 8 based on the performance of the school as a whole on the district exams in each year. All teachers who taught these grades were eligible for the prize. Prizes were awarded in two categories: “Top-scoring schools” and “Most-improved schools.” Schools could not win in more than one category. Improvements were calculated relative to performance in the baseline year. In each category, three first, second, third and fourth prizes were awarded. Out of the 50 schools participating in the program, 24 received prizes of some type, and teachers in most schools should have felt that they had a chance of winning a prize. Prizes ranged in value from 21 to 43 percent of typical teacher monthly salaries.

The comparison of the 50 treatment and 50 control schools suggested that this program did improve performance in the district exams (by about 0.14 standard deviations), but had no effect on teacher attendance. Instead, the teachers held more test preparation sessions. This, the authors conclude, was rational based on the (limited) evidence on what is most effective in improving test scores over the short horizon. However, these preparation sessions are probably poor substitutes for regular classes. This method of pushing up test scores did little for long-term learning, as evidenced by the fact that once the program ended, students who had been in the program schools did not outperform those in control schools.

The lesson from this experiment seems to be that if we want to boost teacher attendance by providing incentives for teachers, the incentive must be tied directly to attendance. Incentives help teachers get a better sense of what their objectives are and will lead them to focus on the most painless way to achieve this particular objective.

Beneficiary Control over Service Providers

The camera example shows that a simple program clearly linking payment to attendance does lead to clear improvement in attendance. When the incentives are focused on test scores, test scores do improve.

An alternative way to improve incentives is to give greater control to the potential beneficiaries. This approach was the main one advocated by the World Bank’s World Development Report 2004 on social services delivery. Shanta Devarajan, who directed the report, summarizes the idea: “Services can work when poor

people stand at the center of service provision—when they can avoid poor providers, while rewarding good providers with their clientele, and when their voices are heard by politicians—that is, when service providers have incentives to serve the poor.”

In order to affect absence, beneficiary control requires two components. First, beneficiaries must have a real demand for the service, so that they feel that they have something at stake in monitoring providers. After all, beneficiary control by its nature faces a collective action problem: the community would like to enforce regular attendance by the provider, but everyone would rather have someone else do the monitoring. It takes some effort for beneficiaries to exercise power, and so for beneficiary control to work, they must feel that this effort is worthwhile.

The second component is that beneficiaries must have a mechanism for affecting providers. In most developing countries, beneficiaries have no way of punishing delinquency. Since the public services are subsidized, the option of beneficiaries voting with their feet is costly, if available (although this alternative is increasingly used—as evidenced by the large number of children in private schools in the poorest Indian states). Having potential beneficiaries seek services elsewhere may also have no impact on absence, since the salaries of providers are rarely determined by usage. In most cases, beneficiaries have no inputs into the hiring or firing decisions, or in determining the actual salaries that get paid. Thus, an array of beneficiary control strategies are possible, ranging from putting the beneficiaries in charge of hiring and firing the providers (or deciding how much they should be paid), to more limited proposals like having the beneficiaries monitor and report provider absences. Making teachers accountable to a school committee or a body of parents is the standard example of this type of reform. These strategies (as the quotation from Shanta Devarajan makes clear) assume that the demand for the services exists and that the only thing beneficiaries are lacking is a way to exert control over the providers.

The advantages of beneficiary control are twofold. First, it may be cheaper for the beneficiaries to monitor the providers: the beneficiaries are both better informed about shirking than whoever was previously responsible for giving incentives (at some higher level arm of the state), and beneficiaries may have means to punish the agent that are not available to others, such as social opprobrium. Second, to the extent that the service is valuable to beneficiaries, they should care more about it and therefore be more willing to reward or punish the agent. In contrast, if after being given control the beneficiaries choose not to enforce the incentives, it must reveal a low demand for the service, which may suggest that trying to enforce a low absence rate was not an especially worthwhile idea in the first place.

Of course, there is no guarantee that beneficiary control will work even if there is demand for the service. In many developing countries, the beneficiaries of education and health services are likely to be socially inferior to the teacher or health care worker, and a government worker may have some power to retaliate

against them. Moreover, in many situations, beneficiary groups may be captured by the service provider through his or her social connections.

Despite the enthusiasm for beneficiary control, few programs that have given communities increased power over providers have been subjected to a randomized evaluation.³ The two randomized evaluations of which we are aware, one in India and one in Kenya, are not encouraging.

Local Monitoring

The first experiment tried to solve the collective action problem of beneficiary control and ensure that monitoring would take place. It was implemented in government health clinics in Udaipur district, Rajasthan (Banerjee, Deaton and Duflo, 2004b), where a member of the community was paid to check once a week, on unannounced days, whether the auxiliary nurse-midwife assigned to the health subcenter was present in the center, and if she was not there, whether she could be found in the village. A parallel system (a monthly visit by a member of the survey team, on the same day) confirmed that this system of local monitoring was properly implemented: external monitors and community members found similar absence rates. However, no attempt was made to impose an external reward system for the nurse-midwives based on the monitoring information.

The idea behind this experiment was to let the villagers choose how they would use the monitoring information they were generating. For example, the villagers could choose to put explicit pressure on the nurse-midwife or try to shame her by exposing her absences. They could even promise her some explicit rewards. In this case, the collective action problem in monitoring stood in the way of effective local control.

The weekly local monitoring system was put in place in 143 randomly selected clinics for eight months. Then, for the next four months, attendance was measured by external monitors carrying out monthly checks in a randomly chosen sample of 80 comparison health centers drawn from the same population from which the treatment centers were previously drawn. Attendance was also measured by external monitors in the treatment centers in each of these four months, while the weekly local monitoring of the treatment centers continued. During those four months, the absence rates turned out to be almost exactly the same in the program and in the comparison facilities (44 and 42 percent, respectively). Community monitoring by itself, even when it is structured and regular, is clearly not enough.

Community Participation

One possible explanation for the lack of impact of community monitoring is that everybody—including the service providers—knows that the community lacks any formal authority to reward or punish providers. In fact, the school committees set up in many developing countries often have no direct say on hiring, firing or pay

³ The EDUCO program, the main example of community management of schools implemented in El Salvador is perceived to be a great success, but the current evaluation does not attempt to correct for the potentially endogenous program placement (Jimenez and Sawada, 2003).

of the teachers. Even when local government bodies are ostensibly given the responsibility of monitoring the health providers and the teachers, they may have no control over pay or postings, which are determined at a regional or national level.

A project implemented by ICS Africa in Kenya tried to address this gap in two ways. First, ICS Africa facilitated a meeting between the school committee and the school administration at the subdistrict level to ensure that the information the school committee had on the functioning of the school was transmitted to the hierarchy, who could then act on it. School committee members were parents from the school, in charge of raising supplementary funds for the school and, in principle at least, of monitoring its functioning. Second, ICS Africa financed a prize that the school committee members could allocate to the teachers who, in their view, performed the best. The prizes (a bicycle for two winners per school, and a set of cutlery for two runners-up per school) should have been significant enough to warrant effort by the teachers. The school committees received some guidance on how to evaluate a teacher's performance, including test scores, attendance and punctuality, and pedagogical methods, but were ultimately left to decide which teachers they wanted to reward.

While the experiment is ongoing, the preliminary results are once again disappointing (Kremer and Vermeersch, 2005). The program was implemented in 36 schools, randomly selected from 72. After one year, the absence rates in treatment and comparison schools were statistically indistinguishable. The children's performance was not any better in treatment schools than in comparison schools. Providing voice and a certain amount of control over resources to the school committees was not enough to reduce absenteeism.

Why Has Beneficiary Control Proved Disappointing?

While there are lots of other ways to motivate and to institutionalize beneficiary control, the experience to date with such programs has proven disappointing. Similar results from Olken (2004) show that increasing participation in community meetings does not result in lower levels of corruption in local development projects in Indonesia, but monitoring by external agencies does.

One possible reason why the results have been the way they have been is that the amount of beneficiary control in the experiments has been quite limited. It is possible that the beneficiaries would have been more enthusiastic about exercising their powers if they had had the power to fire or to select the person to be hired. Supporters of the Education Guarantee Scheme in Madhya Pradesh in India, which gave villagers control over hiring and firing the local teacher, claim that it substantially reduced absence. But given that the scheme also changed the identity of the teacher (it had to be a local person), it is not easy to be sure that it was local control that did the trick. To the best of our knowledge, there has never been a proper randomized evaluation of any scheme of this type.

Another reason for the weakness of beneficiary control is that most communities are not actually particularly upset about the state of education and health

services, even when, objectively, the situation looks dismal. For example, a survey in a rural Udaipur district found that while actual visits to health facilities, as well as some objective health measures, *were* correlated with the quality of the public health facilities, self-reported measures of health and wellbeing, as well as the number of symptoms reported in the last month, were not (Banerjee, Deaton and Duflo, 2004b). Villagers were roughly equally satisfied with their visits to private and public health facilities, and the overall level of satisfaction was high: 81 percent reported that their last visit to a private facility made them feel better, and 75 percent reported that their last visit to a public facility made them feel better. It is true that these reactions were based on visits where the villagers actually found a provider present, but these reactions still seem surprisingly positive given the level of care they were getting (long waits, lack of drugs, having to pay even when care was supposed to be free, almost half the private “doctors” having no medical qualifications).

In short, people in these countries seem to have low expectations from the health care system, and as a result, have little desire to invest time and energy into making it better. When asked about their last visit to a public facility, villagers did not even mention if the facility was actually closed when they arrived: at their actual visit, they always got to see someone. However, despite the higher cost of visiting private facilities, even the poorest households visited public facilities less than a quarter of the time. The rest of the time, they visited traditional healers and *bhopas*. Many households appear to have largely given up on the public sector. In such a case, beneficiary control cannot be the primary tool for fighting absence.

A similar problem arises in education. Recall the experiment in which headmasters were asked to monitor absent teachers, but failed to do so. If the headmaster had enforced the rule of presence or had failed to give the prize to absent teachers, he would have done something that would surely have displeased the teachers, but it seems likely that many parents might not have noticed or cared that efforts to reduce absence were occurring.

If the reason for high absence rates by providers is the lack of demand, then increasing demand for the service may lead to reduced absence rates. Of course, this strategy only makes sense if the presumption is that the demand is actually inefficiently low. This situation could arise because there are externalities across students, or because it is believed that parents or children are not sufficiently attuned to the value of schooling, or because an educated population is a public good. Regardless of whether increasing demand is desirable in itself, the impact of interventions meant to improve demand can offer some insight into absences.

Demand-Side Interventions

There is now a wealth of different ideas for making schools more attractive: improving access by building schools nearer to where people live; improving school infrastructure with investments, such as toilets for girls, lights for rainy days, fans for hot days and sports equipment; improving school inputs, such as textbooks and

charts; and providing incentives for children/parents of children who attend school like direct financial rewards for regularity (such as the PROGRESA program in Mexico) and school meals or indirect rewards such as scholarships for students who do well. Two experiments in Kenya evaluated the impact of demand-side interventions on teacher absence.

Incentives to Learn

An intervention conducted by ICS Africa provided incentives for girls to do well in school (Kremer, Miguel and Thornton, 2004). At the beginning of the school year, ICS Africa announced that it would award scholarships to the highest-scoring 15 percent of grade 6 girls enrolled in the program schools on standardized official district tests, which every child takes at the end of the school year. The program was implemented in a randomly selected half out of 127 schools. The scholarship paid for school fees for the next two years (at that time, students had to pay a fee to attend public school, though this fee has since been removed), as well as a cash payment for school supplies, and offered public recognition at an award ceremony.

In the schools with the scholarship program, both children and teacher presence went up, relative to the comparison schools. Teacher presence (measured at random, unannounced visits) was 6.5 percentage points higher in treatment schools than in comparison schools, reducing absence in program schools by about one-third.

One possible reason for this effect is that the teachers' intrinsic motivation was boosted by the sight of a class full of keen students. Teachers may also have enjoyed higher status when their students received the scholarship. Alternatively, parents may have started taking the teacher absences seriously because it meant that their children were less likely to win a scholarship.

The experiment also provides evidence of externalities associated with demand. Even though only girls were eligible, the presence in school of *both* boys and girls increased in program schools. Moreover, test scores of both boys and girls increased, as did the test scores of girls that were initially low-performing and thus stood a very low chance of winning the scholarship. The most likely channel is that all children benefited from the pressure exerted by the families of the children who were directly eligible. After all, having a teacher present more often will benefit all children in a class.

Incentives to Attend

Most programs of incentives for children do not reward academic performance, but rather school enrollment or student presence in school. For example, the PROGRESA program in Mexico provided cash transfers to family whose children were enrolled in schools and who sought preventative health care. The PROGRESA program had significant effects on school enrollment, but not on student attendance (Schultz, 2000). To our knowledge, its effect on teacher absence has not been studied. In contrast, a child benefits from a school meal

program only if the child attends school. Moreover, if the meal is delivered by the teacher, and would not be delivered if nobody was at school to prepare it, then a school meal program could result in increased participation by both students and teachers.

School meals are quite common—most Indian public schools have them now—but their effect on teacher and student attendance had not been subjected to a randomized evaluation. Vermeersch and Kremer (2005) evaluated the impact of a school breakfast program delivered by ICS Africa in Kenyan preschools on student and teacher presence, as well as student learning. Preschool teachers are locally recruited, and their salary is much smaller than that of regular teachers. Teacher absence rates reach about 30 percent. In this case, however, the meals were not served by the teachers: they were prepared by a specially-hired cook and served by the cook under the supervision of a parent. The school meal program resulted in a 30 percent increase in student participation in school on a given day, but the rate of teacher absence was not affected. We suspect that this result occurred because the parents' demand was focused on the school meal, for which the teacher did not necessarily need to be present. More evidence is therefore needed to understand the effect of school meals on student attendance.

Conclusion

It is hard to resist the conclusion that most attempts to boost the presence of teachers and health providers have not been particularly successful. However, it would be a mistake to read too much bad news into the current evidence. The fact that teachers are willing to work under extrinsic incentives, as they did in the camera project in Udaipur, and do seem to respond to them, is certainly very good news. This finding is especially encouraging given that the incentives were hardly extravagant. For each extra day, the teacher was paid 50 rupees (just over \$6 in purchasing power parity terms), which is close to the average daily wage for the teachers in the control group. Perhaps the innate objectivity of the camera-based system has an intrinsic appeal that had been previously underestimated.

It is also worth recognizing that the working conditions faced by the providers in many of the studies reported here are hardly ideal and that this might partly explain why providers do not always respond to the incentives. Indeed, in some cases the job assignments of teachers and health care workers are poorly defined in a way that makes a seemingly elevated rate of absence almost unavoidable. For example, auxiliary nurse-midwives in rural Rajasthan are assigned to a health subcenter in a village, but they often prefer to live in a nearby town that, for example, has better schools for their children. Six days a week, from a town often 20 kilometers or more away, the midwife has to get to the village by 9:00 a.m., using a combination of very limited public transportation and walking: 6 percent of the

subcenters are more than 5 kilometers from a road. She is then expected to visit all the hamlets near her assigned village once a week, which sometimes involves walking or cycling several kilometers when it is 110 degrees or more in the shade. In these circumstances, most people would probably be tempted to stay home at least occasionally, especially given that they know that if they do not show up, only people in that particular hamlet will ever find out. Rationalizing job descriptions may well be the first step toward getting better attendance in some cases.

On the other hand, the fact that teachers did respond quite strongly to incentives suggests that at least for them, improving attendance is within reach. It is therefore rather disappointing that the beneficiary control mechanisms do not appear to work even in those schools. However, even in that case, the problem seems to be a lack of interest among the beneficiaries (either because they do not value what they are getting, or because the control that they are being offered is too limited to be of interest), rather than that the teachers are too powerful to be affected by parents' demands. When the beneficiaries really wanted more school days, as in the scholarship schools, the teachers seem to have been prepared to deliver. The same lack of interest may also be the reason why external control by headmasters did not do anything: the headmaster faced a choice between doing something that might please some of the parents (though many parents may not have noticed or cared), but would surely displease the teachers.

The lack of public interest in provider absence stems, in part, from a suspicion that what is being delivered in the public facilities is not really very useful. This attitude is perhaps best exemplified by the fact that giving the school meals in preschools did not affect teacher attendance. Parents clearly valued the meals enough to bring in their children, but not enough, given that teachers had nothing to do with the meals themselves, to pressure the teachers into coming more often. This observation is not to deny the evidence suggesting that parents do want their children to go to school—it is only raising the possibility that for many parents, this commitment is perhaps driven more by their sense of what society expects rather than a strong demand for education itself. That sentiment, combined with the oft-expressed feeling that the service providers have a hard enough job in any case, makes the public reluctant to put much effort into improving attendance of public service providers.

It is not, however, clear that this public skepticism about the provision of government health and education services is particularly well-founded. In the case of the health centers, for example, we often heard the complaint that in the government health centers the doctors only give pills, while the private doctors give shots (almost 70 percent of all visits to private doctors lead to a shot). The government protocols do indeed recommend pills for most things, but primarily because pills are cheaper and much safer (they avoid the use of contaminated syringes), and for most things, work just as well. Parents are not ready to pay even a small amount for deworming drugs for their children, even though the drug has

a high rate of return, when benefits are valued over the life of the child (Kremer and Miguel, 2004).

In the case of education, as well, it is hard to imagine that most parents or children in these countries really have a good sense of what the private (and social) return to education would be in ten years for people like them or for their children. It is entirely possible the governments (and so many others) are right in pushing these parents to educate their children. In this scenario it is possible to imagine a virtuous cycle. When teachers start teaching better, education starts to be more valuable. Eventually people get wind of that and start demanding more education. At that point beneficiary control is likely to become much more effective, and this in turn will lead to better schools and even more demand and so on. Such a virtuous circle is perhaps even more plausible in the case of health centers, where people do not have to wait ten years to see the results of better treatments.

■ We thank Angus Deaton, Rachel Glennerster, Rema Hanna, Neelima Khethan, Michael Kremer, Sendhil Mullainathan and Christel Vermeesh for their inputs and comments, and Annie Duflo, Greg Fischer, Shehla Imran, Neeraj Negi and Callie Scott for outstanding research assistance. We are grateful to the Center for Health and Wellbeing at Princeton University, the National Institutes of Health and the John D. and Catherine McArthur Foundation for financial support for the Udaipur health facility survey and the John D. and Catherine McArthur Foundation for financial support for the NFE absence project. Duflo thanks the Alfred P. Sloan Foundation for financial support while conducting this research.

References

Banerjee, Abhijit, Angus Deaton and Esther Duflo. 2004a. "Health Care Delivery in Rural Rajasthan." *Economic and Political Weekly*. 39:9, pp. 944–49.

Banerjee, Abhijit, Angus Deaton and Esther Duflo. 2004b. "Wealth, Health, and Health Services in Rural Rajasthan." *American Economic Review*. May, 94:2, pp. 326–30.

Duflo, Esther and Rema Hanna. 2005. "Monitoring Works: Getting Teachers to Come to School." Mimeo, Massachusetts Institute of Technology.

Glewwe, Paul, Nauman Ilias and Michael Kremer. 2003. "Teacher Incentives." Mimeo, Harvard University.

Jimenez, Emmanuel and Yasuyuki Sawada.

1999. "Do Community-Managed Schools Work? An Evaluation of El Salvador's EDUCO Program." Mimeo, DECRG, The World Bank.

Kremer, Michael and Daniel Chen. 2001. "An Interim Report on a Teacher Attendance Incentive Program in Kenya." Mimeo, Harvard University.

Kremer, Michael and Edward Miguel. 2004. "The Illusion of Sustainability." BREAD Working Paper No. 91.

Kremer, Michael and Christel Vermeesh. 2005. "School Committee Empowerment: Preliminary Notes." Mimeo, Harvard University.

Kremer, Michael, Edward Miguel and Rebecca Thornton. 2004. "Incentives to Learn." NBER Working Paper No. 10971.

Olken, Ben. 2004. "Monitoring Corruption: Evidence from a Field Experiment in Indonesia." Mimeo, Harvard University.

Schultz, T. Paul. 2000. "Impact of PROGRESA on School Attendance Rates in the Sampled Population." Working Paper, International Food Policy Research Institute, Washington, D.C.

Thompson, E. P. 1991. *The Making of the English Working Class*. London: Penguin Books.

Vermeersch, Christel and Michael Kremer.

2005. "School Meals, Educational Achievement and School Competition: Evidence from a Randomized Evaluation." Mimeo, Oxford University.

Voth, Hans-Joaquim. 2000. *Time and Work in England 1750–1830*. Oxford: Oxford University Press.

World Bank. 2004. *Making Service Work for Poor People. World Development Report*. Washington, D.C., and Oxford: World Bank and Oxford University Press.