

Policy note

14/0189

April 2014

**International
Growth Centre**

London School of Economics
and Political Science
4th Floor,
32 Lincoln's Inn Fields
Houghton Street
London WC2A 2AE
United Kingdom

IGC International
Growth Centre

Improving Teacher Attendance using a Locally Managed Monitoring Scheme: Evidence from Ugandan Primary Schools

Jacobus Cilliers^a, Ibrahim Kasirye^b, Clare Leaver^c, Pieter Serneels^d
and Andrew Zeitlin^e

A joint project between: Economic Policy Research Centre (EPRC), the Makerere University School of Computing and Informatics Technology, and World Vision in Uganda, Department of Economics and Blavatnik School of Government at Oxford University, Georgetown Public Policy Institute, and University of East Anglia.¹

EXECUTIVE SUMMARY

Teacher absenteeism remains a serious challenge in Uganda, with estimated rates of absence as high as 27 per cent. Evidence from other countries suggests that a monitoring scheme, combined with bonus payments, could reduce absenteeism, and improve education performance. However, it is unclear what form this scheme should take. **Automated** monitoring via cameras or punch cards is expensive and has been subject to sabotage (Banerjee et al. 2007). **Local** monitoring by head teachers is cheaper and has the potential added value of providing useful information to government for planning purposes. However there is risk of collusion,

^a University of Oxford (contact: jacobus.cilliers@gmail.com)

^b EPRC (contact: ikasirye@eprc.or.ug)

^c University of Oxford (contact: clare.leaver@economics.ox.ac.uk)

^d University of Oxford (contact: p.serneels@uea.ac.uk)

^e Georgetown University (contact: az332@georgetown.edu)

¹ **Please note:** this is a 'Rapid Response' policy paper prepared for the International Growth Centre, and is based on preliminary results from the first term of intervention. An academic article, drawing on results from the first full year of intervention, is in preparation and will be available for circulation in 2014.

Directed and Organised by



www.theigc.org

For media or
communications
enquiries, please
email: mail@theigc.org

especially when financial rewards are attached (Chen et al. 2001). Monitoring by beneficiaries, in this case parents could be attractive, although it is an open question whether parents are able and willing to take on this role. This paper reports on a study assessing the efficacy of alternative forms of local monitoring, both as a means to improve teacher attendance, and to collect good quality monitoring reports.

The study is taking place in 180 rural government schools in six districts in Uganda. At the start of the study, 40 schools were randomly assigned to a control group. Of the remaining 140 schools, 90 were randomly allocated to one of four local monitoring designs, which are the focus of this note. These designs differ along two dimensions: (i) the identity of the monitor and (ii) the stakes attached to the report. In some schools, head teachers have been asked to undertake the monitoring and reporting, whereas in other schools, monitoring and reporting are the responsibility of parents on the school management committee. At the start of the study, all of the nominated monitors were trained to report teacher attendance via SMS on a mobile phone, using a platform developed and managed by Makerere School of Computing and Informatics Technology. The stakes attached to the report vary as follows. In some schools a monthly bonus of US\$ 60,000 (about USD23, or 30 per cent of a monthly salary) is paid to any teacher marked as present in every 'qualifying report' that month, whereas in other schools there is no financial bonus attached and qualifying reports are simply collated and re-broadcast via SMS. For the head teacher's monitoring report to qualify s/he must report every day, while parents need only report once per week. Head teachers thus face a stricter monitoring protocol than parents.

Combining these treatments we obtain four local monitoring designs which are assessed against two criteria: cost effectiveness in inducing higher teacher attendance, and quality of reporting, both in terms of the frequency and reliability of reports. The key findings are as follows:

First, local monitoring improves teacher attendance but only when the head teacher is responsible for monitoring and there are financial incentives for teachers at stake.

Under this scheme, teacher presence is 11 percentage points higher than in the control group. Importantly, this design is also the most cost-effective since bonus payments were no larger than when parents were responsible for monitoring.

Second, although monitors report frequently they all understate teacher absenteeism and parents more so than head teachers.

Head teachers report an average of 2.5 times per week and parents, encouraged to select a single reporting day, report on average once per week. The understatement of teacher absenteeism is substantial across all four schemes: the actual presence rate, as measured by independent surprise visits, is 14 percentage points lower than the reported presence rate. Parents generate significantly less reliable reports than head teachers. There is no evidence that reporting is less reliable when bonuses are paid. Further tests suggest that report quality is wanting because all monitors have a tendency to falsely report absent teachers as present. This is

exacerbated under the parent led schemes because parents opt to report on days when more teachers are present. The stricter reporting protocol under the head teacher led scheme discourages head teachers from doing the same.

In sum, local monitoring schemes **can** improve teacher presence at relatively low cost. However, report quality may be a challenge, with reports systematically understating teacher absenteeism. Intriguingly, although parent monitors were unable to induce higher teacher presence when monitoring alone, preliminary evidence from a further pilot scheme suggests that they can play an important role by auditing the monitoring activity of head teachers. Teachers respond similarly to head teacher led monitoring but parent oversight cuts down on the frequency of false reports. To further reduce absenteeism, the potential benefits of "multiple monitors" are being explored in greater detail in ongoing work.

Taken together, the results summarised in this note underline the intricacy of local accountability systems and point to competing incentives not just for teachers but also for monitors themselves. Local stakeholders care about school performance but have to balance this against questions of morale and the legitimate challenges that face teachers and, of course, against the time and effort required for effective oversight and dialogue. Localized accountability systems are complex, and we are just beginning to learn how technological advances should be matched with the human challenges facing teachers, managers, and parents.

MOTIVATION

Since the adoption of Universal Primary Education in 1997, Uganda has experienced remarkable progress. The number of children enrolled increased from 3 million to 5 million in the first year alone. However, despite a large increase in spending in education over this period, education performance continues to lag (Byamugisha & Ssenabulya 2005).

One possible reason for persistent low education outcomes is teacher absenteeism. A recent study (Chaudhury et al. 2006) estimates that absenteeism rates are as high as 27 per cent in Uganda, the highest observed across countries. Unfortunately, government rarely sanctions absent teachers. This is particularly true in rural areas, where transportation costs and limited inspection budgets often hinder active oversight by district inspectorate teams. Yet, it is precisely in these areas where education challenges are most acute.

Evidence from India suggests that automated monitoring, combined with bonus payments, can reduce absenteeism and improve education performance (Duflo et al. 2012). There are notable downsides, however. First, the Indian program relied on time-stamped cameras placed in each classroom and was expensive to implement. Second, efforts to undertake a similar automated system in the health sector were undermined by both technical and political constraints: the equipment was broken and not quickly replaced, and absences were often deemed exempt (Banerjee et al. 2007).

An alternative is **local monitoring**, whereby school stakeholders –head teachers, parents, or the school management committee– are invited to monitor teacher attendance. An obvious advantage of enlisting head teachers is the low cost associated with their monitoring, as information on teacher presence is already available in the form of daily enrolment registers. Set against this, there is a risk that head teachers falsely report teachers as present, for instance because they are reluctant to withhold bonus payments from absent teachers. Indeed, a Kenyan program that rewarded teachers based on reports submitted by their principal did not have any impact on teacher behaviour (Chen et al. 2001). A potential advantage of enlisting parents as monitors is that, motivated by an intrinsic concern for their children’s education, they may have a stronger desire to punish absent teachers by submitting truthful reports. On the other hand, downsides include higher monitoring costs that may prevent parents from visiting the school on a regular basis and the possibility that parents’ intrinsic motivation may be undermined by fears of reprisals from teachers. Furthermore, parents will only be effective if they truly believe that teacher absenteeism is a problem (Banerjee and Duflo 2009). The efficacy of such local monitoring is therefore an open question.

Beyond the direct interest in improving attendance, local monitoring may also have additional value for planning purposes. Again, there are potential advantages but risks of misreporting. School stakeholders may have information that could prove useful to central government for resource allocation and management decisions. It is difficult for government officials to reach remote schools, yet there is almost always phone reception enabling local monitors to send SMS reports back to central office. But even in the absence of bonus payments, monitors may misreport in order to secure future resources or to reduce the risk of retrenchments. School enrolment information affects school budgets, for example. It is therefore important to determine which type of local monitoring generates the best quality reports.

These two broad questions –the efficacy of local monitoring as a means to improve teacher attendance, and to collect quality information for planning purposes– are particularly policy relevant in Uganda today. The Ugandan government’s Education Management Information System (EMIS) already collects decentralised data and there is donor supported interest to scale this up via the use of mobile phone technology. There is also ongoing discussion about making payments contingent on local reports.

In this study we ask whether local monitoring can play a role in an effective accountability system, and also if it can provide high quality data. To answer these questions, we pilot alternative designs for a mobile-phone-based monitoring system via a randomized controlled trial. We consider the effect of the monitor’s identity (head teacher or parent) and of the stakes involved (bonus payments or information only) on both teachers’ behaviour and on the quality monitoring reports.

STUDY DESIGN

We are carrying out the study in 180 rural government schools in six districts in Uganda. Working with World Vision, we trained school stakeholders (the head teacher and/or parents) how to report teacher attendance via SMS on a mobile phone, using a platform developed and managed by Makerere School of Computing and Informatics Technology. These reports are sent to a central database in Makerere University, aggregated to school-level, and a summary of the results is re-broadcast back to school stakeholders via SMS.

At the start of the study, we randomly assigned 90 schools to one of four basic reporting schemes and 40 schools to the control group, where no reporting took place.² The reporting schemes differ along two dimensions. First, in some schools, head teachers have been asked to undertake the monitoring and reporting, whereas in other schools, this is the responsibility of parents on the school management committee (SMC). Second, in some of the schools we pay a monthly bonus of US\$ 60,000 (about USD23, or 30 per cent of a monthly salary) if a teacher is marked as present in each of the weekly ‘qualifying report’ during that month. Qualifying reports are selected according to the following procedure. For the schemes where a parent monitors, we randomly select one of the reports submitted that week. For schemes where the head teacher monitors, we randomly choose one day that week and then a report (if there is one that day). This means that head teachers must monitor and report every day to ensure that bonuses are paid out, whereas parents need only monitor and report once per week.³

In summary, schools were allocated to one of five groups:

- No-one reports (Control): 40 schools
- Head teacher, information only: 20 schools
- Head teacher, bonus payments: 25 schools
- One parent on SMC, information only: 20 schools
- One parent on SMC, bonus payments: 25 schools.

The focus of this note is on the impact of head teacher and parent monitoring with and without bonus payments (schemes one to four). These interventions were launched in September 2012, with training by World Vision.

Independently of this intervention, we collected our own data on teacher presence in these schools through unannounced ‘spot checks’. The first such measurement exercise, a baseline, was undertaken in July 2012. We revisited the schools and conducted further unannounced spot checks during November 2012, April/May 2013, and August 2013. The present note focuses on results from the first term of implementation, drawing primarily on data from the November spot checks and intervention-generated reports from the corresponding period.

² The 50 remaining schools were assigned to a pilot of more complex reporting schemes. Under these schemes, a teacher can only qualify for a bonus payment if **multiple** monitors mark him/her present on the same day. Preliminary results from this pilot are discussed in the section on multiple monitors.

³ We designed the protocol in this way in anticipation of the difficulties parents were likely to face in monitoring several times per week.

FINDINGS


As discussed above, we want to evaluate the schemes against two criteria: (i) cost-effectiveness in inducing higher teacher presence and (ii) quality of monitoring reports, both in terms of the frequency and reliability of these reports. The results presented are based on the first term of implementation, ending in November 2012. Further results are the subject of on-going analysis.

Presence: head teacher LED monitoring with bonus payment substantially improves teacher presence

First, we measure the impact of each reporting scheme on teacher presence. In Figure 1, we plot the proportion of all independent spot checks where a teacher was present across the different reporting schemes. The red line shows the presence rate in the control group, where no reporting took place. The first two bars indicate schools where the head teacher is asked to monitor and report, while the final two bars indicate schools where a parent on the school management committee is asked to monitor and report. The first and third bars indicate schools where reports are for information only, while the second and fourth bars indicate schools where reports can also trigger bonus payments for teachers.

Head teacher led monitoring with bonus payment is the only scheme that induces a statistically significant increase in teacher attendance relative to the control group. In this scheme, teacher presence is 73 per cent, 11 percentage points higher than in the control group. By contrast, effects of schemes without bonus payments, and schemes managed by parents alone, have more muted, and statistically insignificant, effects.

Figure 1: Presence rates across different reporting schemes.




Note: Figure based on 2181 teacher-days with independent spot checks in November.

Next, we compare the direct cost associated with bonus payments. The dark red bars in Figure 2 show the proportion of weeks with a qualifying report that marks the teacher as present. This proportion is lower under the scheme managed by a head teacher than under the scheme managed by a parent, although the difference is not statistically significant. It is thus more effective but no more expensive when head teachers are asked to monitor and report.

It is worth noting that the bonus payout rate is low under the head teacher led scheme because head teachers often fail to comply with their stricter reporting protocol. (Recall that head teachers had to submit a report every day to be certain of triggering a bonus payment, whereas parents only had to submit once a week.) The light red bars in Figure 2 show the proportion of weeks with any qualifying report. The proportion of weeks where a head teacher submitted a report (on the day that we randomly selected) is also lower than the proportion of weeks where a parent submitted a report (on any day that week).

Figure 2: Cost associated with bonus payment across schemes.



Note: Figure based on 3436 teacher-weeks in November across two basic schemes with bonuses.

We conclude that parent monitors in the above schemes are not effective in improving teacher presence, while head teachers are only effective when bonuses for teacher presence are involved. Next we investigate which schemes are more effective at generating good quality monitoring reports.


QUALITY OF REPORTS

We focus on two key dimensions of reporting quality: frequency of reporting and reliability of reports, and discuss each in turn. Observing gaps in reliability, we then explore the potential reasons, distinguishing between false reporting and differences in the selection of reporting days.

Frequency: Reporting is more frequent when linked to bonus payments

Figure 3 shows the total number of reports per week across the different schemes. Total frequency is high: on average 1.9 reports are submitted per week. Unsurprisingly, given that their intervention protocol called for daily rather than weekly reporting, head teachers report far more frequently than parents: on average 2.5 times per week compared to once per week. Furthermore, both parents and head teachers make more reports when there are bonus payments available: on average, there are 0.5 more reports per week when bonuses are attached, a difference that is statistically significant at the 5 per cent level.

Figure 3: Total reports per week.




Note: Figure based on 6288 teacher-weeks in November across all basic schemes.

Reliability: All monitors over-state teacher presence but parents far more so

Next, to determine if monitors' reports are a true reflection of actual presence rates, we restrict the sample to all weeks where there is both an independent spot check and a qualifying report from the head teacher/parent (the rectangle in Figure 4). We then compare the actual presence rate (based on spot check days in the smaller ellipse) with the reported presence rate (based on all reporting days in the larger ellipse) during those weeks.

Figure 4: Spot check and reporting days in weeks with both a spot check and qualifying report.


Days in weeks with both spot check and report (9,996 days, 1428 weeks)



In Panel A of Figure 5 the (first) blue bars plot the actual presence rate and the (second) red bars plot the reported presence rate. Panel B plots the difference.

Figure 5: Actual versus reported presence rates.


Panel A – actual and reported presence



Note: Figure based on 1,428 teacher-weeks with 1 selected report and >=1 spot check per week.

Figure 5: Actual versus reported presence rates.

Panel B – difference between actual and reported presence



Note: Figure based on 1,428 teacher-weeks with 1 selected report and >=1 spot check per week.


Figure 5 shows that there is a large discrepancy between actual and reported presence. On average, the true presence rate is 14 percentage points **lower** than the reported presence rate. Somewhat surprisingly, this discrepancy is much larger under the parent-led schemes. The reports from parents overstate actual presence by 18 percentage points, compared to 10 percentage points for the reports from head teachers. This difference is statistically significant at the 10 per cent level. There is no evidence that the discrepancy is larger when bonus payments are available; it is actually slightly smaller, although this difference is not statistically significant. To explore the possible causes for the discrepancy between actual and reported presence, we distinguish between two forces: false reporting and selection.

False reporting is considerable but similar across head teachers and parents


The most obvious explanation for the discrepancy between actual and reported teacher presence plotted in Panel B of Figure 5 is that monitors falsely report absent teachers as present. To test for this, we restrict the sample to all days where there was a simultaneous spot check and report by a monitor (the intersection of the two ellipses in Figure 5) and compare actual and reported teacher presence (Figure 6). In Panel A of Figure 6 the (bottom) darker blue area in the bars indicates the proportion of teachers that are both present and reported as present, reflecting truthful reporting of teacher presence. In the (second from bottom) light blue area, teachers are present but falsely reported as absent. The dark red area indicates the proportion of teachers that are absent but falsely reported as present. In the (top) light red area, teachers are absent and also reported as absent, reflecting truthful reporting of teacher absence.

Figure 6: The extent of false reporting.

Panel A – actual and reported presence



Panel B – difference between actual and reported presence



Note: Figure based on 998 teacher-days with reports and independent spot checks.

Panel A of Figure 6 indicates that there is false reporting under all four local monitoring schemes, with between 11 and 13 per cent of teachers falsely reported as present. As one might expect, far fewer teachers are falsely reported as absent, just 1 to 3 per cent. Subtracting the latter from the former –that is, the light blue bar from the red bar– we obtain the quantity plotted in Panel B: the difference between actual and reported presence for the subsample of days with both a spot check and report. Notice that there is no variation in false reporting between the parent led and head teacher led schemes. Even though on average over weeks with both a spot check and qualifying report parents over-state presence


by more than head teachers, they are no more likely to submit false reports. It follows that false reporting can only explain part of the pattern in Figure 5. To explain this puzzle, we turn to the second explanation: selection.

Parents select reporting days with higher teacher presence

Another reason for the discrepancy between actual and reported presence in Figure 5 could be variation in teacher attendance and reporting **within** weeks. To illustrate, suppose that teacher presence is high on Mondays and, further, that monitors are more likely to report on Mondays than on other days. Reported presence will (mechanically) be higher than actual presence because the latter is based on a random sample of days whereas the former is based on a ‘selected’ sample of days when teacher presence is high. Of course, there are many reasons why both teacher attendance and monitoring might be higher on some days than others. Both teachers and monitors might be absent on market days, for example; or maybe they see virtue in coordinating their efforts and actively choose to attend and report on the same days. A key question is whether such variation within weeks is likely to be higher in schools where parents are asked to monitor than in schools where the head teacher is asked to monitor. Here the answer is unequivocally yes: head teachers face a protocol that actively discourages the ‘selection’ of a particular day for attendance and reporting, whereas parents do not.

To test for such selection, we now restrict the sample to the days where independent spot checks took place (the smaller ellipse in Figure 4). These are the days that were used to plot the dark blue bars in Panel A of Figure 5. We then disaggregate by whether the spot check day happens to coincide with a day on which the monitor chose to report, or not. In Figure 7, the (first) blue bars plot the actual presence rate on the days that do not coincide with days that the monitor submitted a report (blue region of the smaller ellipse in Figure 4). The (second) green bars plot the actual presence rate on the days that do coincide with days that the monitor submitted a report (green region, intersection of the two ellipses in Figure 4).⁴ Comparing these bars, we can test if parents reported on days where presence rates are higher.

Figure 7: Teacher presence on days with and without reports.



Note: Figure based on 1,652 teacher-days with spot checks in weeks with both check and qual reports.

Figure 7 shows that this is indeed the case. Under the parent led schemes, the actual presence rate is 12 percentage points higher on days when a report is submitted than on days when a report is not submitted. In contrast, under the head teacher led schemes, the actual presence rate is slightly **lower** on reporting days. The difference between the parent led and head teacher led schemes is significant at the 5 per cent level. Thus selection explains the puzzle noted above. On average over weeks with both a spot check and qualifying report, parents over-state presence by more than head teachers, not because they are more likely to submit false reports, but because they are more likely to monitor on selected days (presumably as a consequence of the more flexible reporting protocol).

To sum up, the quality of reports is poor under all of the locally managed monitoring schemes due to false reporting and is **worse** in the parent led schemes due to reporting on selected days. Revisiting the earlier results it now seems remarkable that, **in spite** of false reporting, the head teacher led scheme with bonus payments is nonetheless able to induce higher teacher presence.

WAYS FORWARD: MULTIPLE MONITORS, WITH PARENTS AS AUDITORS

Results so far do not seem to favour parents as monitors. However, findings from another reporting scheme do suggest an important and intriguing role for parents in improving outcomes. While the day-to-day burden of regular monitoring activities appears to be a barrier to effective parent led monitoring schemes, our results suggest that use of **multiple monitors** – both head teachers and parents – can achieve attendance gains similar to the head teacher managed scheme, at lower cost.

⁴ Note these are the days used to construct Figure 6.

Specifically, we are piloting a scheme in which head teachers carry the primary burden of monitoring – submitting daily attendance logs – and parents play the role of auditors. A teacher can only qualify for a bonus payment if both the head teacher and parent mark him/her present on the same day. What we see so far suggests that under this scheme, teachers respond similarly to head teacher led monitoring, but parent oversight cuts down on the frequency of false reports, thereby limiting the payment of bonuses in cases where this was not merited by teacher behaviour. It appears that the most cost-effective way to improve teacher attendance is thus to make use of both head teachers and parents in a reporting scheme.

CONCLUSION

In this ongoing study, we are assessing different locally managed monitoring schemes. We have experimentally varied the monitor's identity (head teacher or parent) and the stakes involved (bonus payments or information only) and have evaluated these schemes in terms of their ability to induce higher teacher presence at a low cost, and also to produce frequent and reliable monitoring reports.

Our results indicate that local monitoring can increase teacher presence, but only under a head teacher led scheme where bonus payments are available. This is also the most cost-effective of the four schemes discussed in this note: the bonus payout is no larger than under the parent led scheme with bonuses attached.

Turning to the quality of monitoring reports, frequency of reporting is high across all schemes, but far higher under the head teacher led schemes. Reporting is also more frequent under the schemes where bonus payments are available.

A worrying finding is that **all** of the locally managed monitoring schemes generate data that systematically understate teacher absenteeism, and far more so in the parent led schemes. In weeks with both an unannounced spot check and a qualifying report, the reports from parents over-state actual presence by 18 percentage points, compared to 10 percentage points for the reports from head teachers. Perhaps surprisingly, the discrepancy is similar when bonuses are attached, suggesting that there is not a trade-off between increasing the stakes of monitoring and quality of reports. Further tests reveal the mechanisms at work. False reporting is high across all schemes (10 per cent of the reports submitted on days with an unannounced spot check falsely mark an absent teacher as present) but parents are no more likely to submit a false report than head teachers. Parents are, however, more likely to engage in selective reporting. Teacher presence is higher on days when parents choose to report than on days when they choose not to report. This is not the case for head teachers (presumably) because the monitoring protocol actively discourages selection. Although pilot interventions to date do not allow a definitive explanation, this suggests a potential lesson for the design of local monitoring schemes: since interested parties may find it easier to report selectively than to lie outright, stricter protocols to discourage this type of selective reporting may improve the effectiveness of even parent led monitoring schemes.

Our study also points to an intriguing avenue for the further improvement of local monitoring schemes. Although parent monitors were unable to induce higher teacher presence, evidence from a further pilot scheme suggests that they may be able to play an important role in decreasing the cost of an accountability system by auditing the monitoring activity of head teachers. We are exploring the potential benefits of “multiple monitors” in greater detail in on going work.

Finally, it is worth emphasising that the impacts of the schemes reported here may change over time as both teachers and monitors learn and adapt. Attendance might increase as teachers place greater trust in the system having received their first bonus payments. Alternatively, false reporting might become more prevalent as absent teachers take more active steps to ensure bonus payments. We plan to analyse these issues using data collected during 2013. Localized accountability systems are complex, and we are just beginning to learn how technological advances should be matched with the human challenges facing teachers, managers, and parents.

REFERENCES

- Banerjee et al. 2007 Banerjee, A., Duflo, E. & Glennerster, R. (2007), 'Putting a band-aid on a corpse: incentives for nurses in the Indian public health care system', *Journal of the European Economic Association* **6**, 487–500.
- Banerjee, A., Duflo, E., (2009), 'Addressing Absence', *Journal of Economic Perspectives* **20**(1), 117–132.
- Byamugisha & Ssenabulya 2005 Byamugisha, A. & Ssenabulya, F. (2005), 'The SACMEQ II project in Uganda: A study of the conditions of schooling and the quality of education', *Harare, Zimbabwe: SACMEQ* .
- Chaudhury et al. 2006 Chaudhury, N., Hammer, J., Kremer, M., Muralidharan, K. & Rogers, F. H. (2006), 'Missing in action: Teacher and health worker absence in developing countries', *Journal of Economic Perspectives* **20**(1), 91–116.
- Chen et al. 2001 Chen, D., Glewwe, P., Kremer, M. & Moulin, S. (2001), 'Interim report on a preschool intervention program in kenya', Mimeo, Harvard University.
- Duflo et al. 2012 Duflo, E., Hanna, R. & Ryan, S. (2012), 'Incentives work: Getting teachers to come to school', *American Economic Review* **102**(4), 1241–1278.