

Teachers Pay and Educational Outcomes: Evidence from the Rural Hardship Allowance in Zambia¹

Grieve Chelwa
Graduate School of Business
University of Cape Town

Center for African Studies
Harvard University

Miquel Pellicer
University Duisburg-Essen

SALDRU
University of Cape Town

Mashekwa Maboshe
School of Economics
University of Cape Town

¹ We would like to thank the following people: Soren Henn for sharing his data on GPS coordinates for district centers, Benjamin Chibuye for immeasurable support during the writing of this paper, Bupe Musonda of the Ministry of Education for availing the Annual School Census data to us, Michael Chilala for granting us access to the Grade 7 Examinations data and Grevazio Kapanda for able research assistance. We acknowledge funding from the International Growth Center (IGC).

Abstract:

This paper studies the effect of unconditional teacher salary increases on teacher and student outcomes. To study the issue, we evaluate the rural hardship for teachers in Zambia, which corresponds to salary increase of 20%. This allowance is allocated to schools beyond a given distance from their district center, and this allows us to use a regression discontinuity design to identify its effects at the school level. We use administrative data from 2004 to 2015 on school and teacher characteristics and on Grade 7 student test scores. In addition, we perform a small telephone survey of schools close to the eligibility threshold for the allowance. We find that crossing the eligibility threshold increases the share of teachers obtaining the allowance by 40pp. However, we find no significant effects on teacher attrition, teacher characteristics or student grades. Our null results on student performance is inline with recent research on the topic. The null results on teacher attrition are more surprising, and we discuss possible reasons that may explain them. Most notably, using our telephone survey, we provide suggestive evidence that teacher departure decisions are more based on distance to amenities and to delays in payment than on the receipt of the allowance.

1. Introduction

There has been an explosion of research in recent years on the factors that can improve student learning in developing countries (see recent reviews in Snijtsveit et al. 2016 and Glewwe et al. 2011). Teachers are one type of input that has received attention. Evidence from high income countries shows that teachers can have a large role in student learning and long term outcomes (Jackson et al. 2014; Chetty et al. 2014), and recent evidence from Pakistan suggests a high impact of teachers on grades in a developing country setting as well (Das et al. 2014).

There have been several studies evaluating interventions aimed at increasing the productivity of teachers in developing countries. Most of these interventions have explicitly provided “hard” incentives to affect specific types of teacher behaviors. For instance, teacher presence has been incentivized by monitoring or providing financial incentives conditional on presence (Duflo et al. 2012), and teachers contribution to student performance has been incentivized by providing financial rewards to teachers conditional on positive student test scores (Muralidharan et al 2011, Kremer et al. 2010). While some of these interventions have proven successful, a recent

review comparing the effect of different types of interventions on student learning concludes that the effects of teacher incentive interventions have been generally small (Snilstsveit et al. 2016). Moreover, some studies find that such interventions have generated undesirable outcomes (such as teaching-to-the-test, see Kremer et al. 2010).

There has been growing interest in the potential role of “soft” approaches to increase teacher productivity, that tap into behavioral responses such as reciprocity, or intrinsic motivation (Besley and Gathak 2014). In particular, higher *unconditional* wages may improve the productivity of public servants via selection of more motivated teachers (Dal Bo et al 2013), or via reciprocity.

This paper studies the effect of unconditional increases in salaries on teacher and student outcomes. We do so by evaluating the effects of the rural hardship allowance in Zambia, which provides a 20% increase in salaries paid to teachers in rural schools. The rural hardship allowance is allocated to schools outside a given radius from district centers, and this allows us to estimate the effects of salaries using a regression discontinuity approach. We study the effect of the allowance on teacher attrition, teacher characteristics, and student grades.

Several recent studies study the effect of unconditional wage increases on teacher and student outcomes in developing countries: in Brazil, Uruguay, Pakistan and Indonesia (Andrade da Silva Filho et al. 2014, Cabrera et al 2016, Das et al. 2016 and De Ree et al. 2015). All these studies tend to find either no effect or a small effect on student performance. However, there are several reasons why more evidence is necessary to bring to bear on the question. First, the studies on Brazil and Pakistan (Andrade da Silva Filho et al. 2014 and Das et al. 2016, respectively) employ either a simple difference or a differences in differences estimator which may not be able to fully account for confounding factors. The studies on Uruguay and Indonesia (Cabrera et al 2016 and De Ree et al. 2015, respectively) address this problem by using experimental and quasi-experimental methods. However, there are still reasons why their results may not generalize easily. Uruguay and Indonesia are middle income countries with very

different contexts from low income countries where teacher salaries are very low and unconditional salary expansions may be most relevant. In addition, the low estimates in De Ree et al. (2015) could be partly due to their specific design. They study a policy where salaries of teachers are doubled, using a randomized phase-in approach where the control group knows they will get the salary but has not gotten it yet. If reciprocity is a relevant channel through which unconditional wages affect teacher performance, teachers in the control group may change their behavior even if their salary increase has not yet materialized.²

This paper contributes to this nascent literature by using a quasi-experimental approach that evaluates an existing policy in Zambia, a relatively low-income country. In addition, our analysis provides a direct evaluation of the rural hardship allowance in Zambia, an important policy that seeks to improve the level and equitable distribution of educational outcomes in the country. Indeed, the rural hardship allowance was designed to reduce the relatively high teacher attrition experienced by schools in rural areas, where educational outcomes tend to be weakest, and thereby improve these outcomes.

We construct a school level dataset merging the Zambian Annual School Census and the Grade 7 Examination results to obtain information on teacher attrition, teacher characteristics and school grades in around 3000 schools from 2004 to 2015. The current allocation rule for the rural hardship allowance was established in 2010, and this implies that we can use outcomes pre- and post-treatment. In addition, we obtained a list of schools actually receiving the hardship allowance in 2017 from the government's payroll department.

The rule assigning the allowance is based on distances computed from GPS coordinates between district centers and the schools. This renders virtually impossible a manipulation of the running variable and lends credibility to our approach. Indeed, balance checks show that pre-treatment outcomes of schools at either side of the cutoff are similar.

² In turn, the specific characteristics of Cabrera et al (2016) may also explain their low estimates. They evaluate a bonus incentive scheme for teachers to relocate to poorer schools; as they mention, the scheme was taken up mostly by experienced teachers, and these may be the type of teachers with low wage elasticities.

We find that the rule is not implemented consistently, leading to a first stage lower than anticipated. There are two reasons for this. First, some schools get reclassified ex-post if the GPS distance computed is considered to be a misleading measure of remoteness of the school.³ This problem reduces our first stage coefficient from 1 (if the rule were implemented perfectly) to around 0.45. Second, there is teacher-payroll mismatch in Zambia (Auditor General, 2014). The payroll database that determines the payment of salaries and allowances sometimes features teachers at schools where they no longer teach. This implies that there can be teachers obtaining the allowance teaching at schools ineligible to get the allowance. Because there is no official information on this, we conducted a telephone survey of head teachers for schools around the eligibility cutoff, stratified in pairs of schools close to each other. We succeeded in obtaining information for 137 schools, corresponding to 44 matched pairs, one at either side of the threshold. Taking this into account, the first stage coefficient drops slightly to around 0.4. Nevertheless, the instrument remains strong with F-statistics greater than 10 in all specifications.

Consistently with the literature above, we find no effect of the rural allowance on student outcomes, suggesting that unconditional salaries have little effect on teacher performance, even in low income countries. More surprisingly perhaps, we find that the allowance has no statistically significant impact on teacher attrition or teacher characteristics. We do find consistently positive effects on the stock of teachers and on teacher tenure, suggesting the possibility that the allowance may have a small contribution to keeping teachers in rural stations, but the evidence for this is weak. We discuss possible explanations for these insignificant effects. First, it is possible that we lack power to identify effects, if these effects are small. The reduction in our first stage implies that the jump in average salary upon crossing the threshold is smaller than 20% (estimated at about 8% from our telephone survey). Second, we consider that the “binding constraints” to keeping teachers in rural areas may be other factors such as infrastructure and distance to amenities rather than salaries. We provide suggestive

³ For instance, we were told that if there is a natural barrier such as a mountain between the school and the district center, schools closer to the center than the threshold might be granted the allowance

evidence from our telephone survey that distance to amenities and delays in payments are more relevant for explaining teacher departures from schools around the threshold than salaries received through the rural allowance. If there are complementarities between salaries and amenities in the teacher's utility function, the allowance could have a stronger effect if these amenities were present.

The paper is organized as follows. Section 2 provides some background on education in Zambia and on the rural hardship allowance for teachers. Section 3 describes the data and provides some descriptive statistics, while section 4 explains our empirical approach. Section 5 shows the results, section 6 discusses them, and section 7 provides some brief concluding remarks.

2. Background: Education in Zambia and the Rural Hardship Allowance

Zambia's education sector faced substantial setbacks following the economic crisis that began in the mid-1970s and the ensuing structural adjustment policies (SAPs) of the 1980s and 1990s. The country's expenditure on education as a percentage of GDP declined from 5% in the 1960s and 1970s to 2% in the 1980s and 1990s (World Bank, 2014). As a result, the pace of school construction and of teacher recruitment slowed down and did not keep up with population growth. Further, the reduced budgetary allocations to the sector served to make teaching a less attractive profession and many teachers left the profession or emigrated to neighbouring countries. Not surprisingly, the teacher-pupil ratio increased from about 40 in the 1970s to 50 in the 1990s and was even as high as 80 in the more rural parts of the country (Government of the Republic of Zambia, 2006). This occurred even when the primary school net enrolment ratio declined from 80% in the 1970s to 70% in the 1990s (ibid). These shocks to the country's education sector are likely behind the less than satisfactory performance on internationally standardized tests. Zambia has, for instance, consistently performed at the low end of the Southern and Eastern Africa Consortium for Monitoring Education Quality (SACMEQ). SACMEQ administers tests to assess the level of reading and maths abilities among Grade 6 pupils in the region.

The government, recognizing the challenges faced by the education sector, instituted a number of policy responses in the last decade. The Basic Education Sub-Sector Investment Plan (BESSIP), which ran from 1999 to 2002, sought to increase access to and the quality of basic education in the country. One of BESSIP's landmark achievements was the 2002 abolition of school attendance fees for Grades 1 to 7. Whereas BESSIP was largely successful in enhancing access, school quality suffered in its wake because it focused less on teacher recruitment and retention. The Ministry of Education's subsequent policy plans (MoESP, NIF II and NIF III)⁴ have thus focused on improving quality primarily through the large-scale recruitment and retention of teachers. For instance, NIF II which ran from 2006 to 2011, set itself the target of recruiting 5000 teachers every year (Government of the Republic of Zambia, 2006). NIF III, set out to recruit 3000 teachers every year over the period 2011 to 2015 (Government of the Republic of Zambia, 2011).

With such increases in teacher numbers came the concern that the quality of instruction might suffer, and the challenge to keep the teaching profession attractive. The government then instituted significant increments in teachers' basic pay over the last decade or so, the most significant of which occurred in 2013 and saw salary increments of up to 200%. In addition to this, a variety of incentive schemes have been devised (such as salary increments, allowances, teachers' houses, training, etc...) with the aim of keeping and motivating teachers.

To this end, the government implemented a rural/remote hardship allowance. The scheme first emerged in the 1990s but was of a small quantity and plagued with problems.⁵ In 2008, a substantial rural hardship allowance corresponding to 20% of the base salary was established for all public servants. The rule governing eligibility was a complex combination of distance of the rural station to various amenities (the rural station - clinic school, etc. - ought to be more

⁴ MoESP stands for Ministry of Education Strategic Plan. It ran from 2003 to 2007. NIF II stands for National Implementation Framework II. It ran from 2006 to 2011 and was the framework guiding the Ministry of Education's implementation of the Fifth National Development Plan. NIF III guided the Ministry of Education's efforts in this regard between 2011 and 2015.

⁵ Personal communication, Ministry of Lands.

than 20 Km from the nearest bank, 10 Km from nearest police station, etc.). In 2010, the rule for teachers changed, and was dramatically simplified. It was decided that the single criterion would be distance to the nearest district center. Schools beyond a pre-specified cut-off would qualify for the rural hardship allowance. Districts were divided into four categories according to their degree of remoteness and the cut-off was set differently for each of these categories. For instance, the most remote districts had a cut-off of zero (so that all schools qualified to obtain the allowance), moderately remote districts had a cut-off of between 20 and 25 Km and the most urban districts had a cut-off of 30 Km.

We met several government officials to verify that this rule was actually implemented to pay the allowance. The rural hardship allowance is paid directly to teachers by the government's payroll department on the bases of their database of schools and of school eligibility. The eligibility of each school is determined by the Ministry of Lands, which collects GPS coordinates of schools and computes distances to the nearest district centers. We verified that the rule used by the Ministry of Lands was indeed the one promulgated in 2010. Schools are allowed to contest their allocation if, for instance, the school is separated from the nearest district center by natural barriers that makes the actual travel distance much longer than the GPS distance. In those cases, the eligibility status can be changed. This implies that the actual receipt of the allowance is not completely determined by the rule.

One feature of the teacher pay system in Zambia which is potentially problematic for our analysis is that there is a mismatch between the schools where teachers are paid and where they actually teach (Auditor General, 2014). The government's payroll department pays salaries (and allowances) on the basis of their database, but it appears that the database is not kept up-to-date. This implies that when teachers move to a new school they may still appear in the payroll as being part of the former school. This also applies to the rural hardship allowance, and this implies that eligible schools may have some teachers that do not receive the allowance and vice versa. The size of the problem is unknown, but there are indications that it could be significant. The Auditor General (2014) provides an example of a school where half of its

teachers were paid elsewhere. This could substantially reduce the strength of our instrument, if the problem is widespread, because it implies that the increase in salaries upon crossing the eligibility threshold is less than 20%. Because this issue has potentially severe implications for our analysis, we decided to undertake a telephone survey of schools around the threshold to ascertain the extent of the problem and its implications for our analysis. As it happens, these implications turn out to be minor.

3. Data

Our empirical analysis is based on two types of data: administrative data on schools including school-level information on allowance receipts and examination grades, and a telephone survey we conducted ourselves with head teachers close to the hardship allowance eligibility threshold.

3.1. Administrative School Data

The main data source for our analysis is the Annual School Census collected annually by the Ministry of General Education (MoGE). The census forms a vital part of the Ministry's annual planning and programming activities. It contains a rich set of data on the characteristics of around 9,000 schools across the country of which about 5,000 are run by the Zambian government. Our analysis is restricted to the government-run schools as these are the only ones for which the rural hardship allowance is applicable.

At the beginning of each year, each school registered with the Ministry of General Education (MoGE) is sent a questionnaire. The questionnaire is completed by the head teacher and returned to the MoGE in the middle of the year. The questionnaire elicits a comprehensive set of information on teachers, pupils and the school itself. Head teachers are asked to fill-in information on the qualifications (professional and academic), tenure, age and gender of teachers. Information collected on the school itself includes the level of the school (basic or high school), entity responsible for the school (government, private, church, etc...), year the school

was established and school infrastructure (desks, blackboards, toilets) among others. The characteristics of pupils (number, age, gender, etc...) are also collected.

Fortunately for our purposes, the census also contains GPS coordinates for over 80% of government-run schools every year. We construct the running variable of distance to the nearest district center ourselves, using the school GPS coordinates and the GPS coordinates of district centers taken from Henn (2016).

The Annual Schools Census has, in principle, been conducted every year since 2000. Unfortunately, there were some inconsistencies in the way variables were defined between 2000 and 2004 making it difficult to use the data from earlier years. Therefore, our analysis only makes use of census data from 2004 to 2015, with the exception of the 2010 census whose files are not on the MoGE database. The absence of the 2010 files does not present much of a problem given that 2010 is, in any case, the year in which treatment starts.

We use the Annual Schools Census to derive the main teacher outcome variables which are the transfer rate, the stock of teachers at a school and the average tenure of teachers at a school. The transfer rate measures the percentage of a school's teachers who transfer out of that school to another school in a given year. Unfortunately, the census does not give additional information on the reasons for transferring out of a school. For example, teachers might transfer to another school to follow a spouse or might be asked to transfer to fill a vacancy at another school. It appears, however, that many transfers from rural schools in Zambia are in the main motivated by hardship concerns (Mulcahy-Dunn, et al. 2003). The other two main teacher variables are straight forward: the stock of teachers measures the number of teachers at a school. The average tenure measures the average number of years that teachers have continually served at a school.

In addition to these, we also investigate two other outcome variables: the average years of schooling and the average age of teachers at a particular school in a given year. Much like the

main outcome variables, these latter variables are constructed from data on individual teacher characteristics contained in the Annual Schools Census.

We use data on outcome variables from before 2010 to check for pre-treatment balance, and as control variables to obtain more precise estimates in our analysis.

The Annual Schools Census does not contain information on whether a school actually gets the rural hardship allowance. We obtained this information from the Payroll Management Establishment Control (PMEC) of the Government of Zambia. PMEC maintains the payroll of all civil servants in Zambia including teachers. The challenge is that the list from PMEC, aside from containing names of schools and districts in which schools are located, does not have numeric identifiers allowing us to seamlessly merge it with the Annual Schools Census. The merging was done manually with the Annual Schools Census using both the name of the school and the district in which the school is located.

3.2. Grade 7 Examinations Data

We also seek to investigate whether the hardship allowance has any impact on learning outcomes via any impact it might have on teachers. Given that the Annual School Census does not contain any information on learning outcomes, we combine it with data on school-level performance on the Grade 7 Examinations. The Grade 7 Examinations, sometimes referred to as the Primary School Leaving Examinations, are the first high stakes exams in the Zambian school system. Performance on these exams determines whether candidates proceed onto secondary school. We have school-level performance data on the Grade 7 exams from 2010 to 2014 provided by the Examinations Council of Zambia (ECZ). Individual performance on the exam is classified into one of four categories: Division One, Division Two, Division Three and Division Four. Division One is the highest level of achievement with Division Four being the lowest. We use the percentage of students at a school who score a Division One as the main outcome measure. The analysis is done separately for boys and girls.

Even after merging the Annual Schools Census with the examinations data, we end up with a situation whereby not all schools in the Census have associated examinations data. Two reasons account for this. The first is that not every school is an examinations center as schools have to be assessed as having the prerequisite infrastructure to hold exams. Schools that have candidates but are not examinations centers are required to send their candidates to the nearest qualifying school. The second reason is specific to the Annual Schools Census and the Grade 7 Examinations datasets. Whereas both datasets have unique school identifiers allowing the identification of schools across time, there is no unique identifier allowing us to directly link schools between the two datasets. However, in 2008 an attempt was made by the ECZ to link the EMIS number (Annual Schools Census unique school identifier) with the ECZ Facility Code (the unique school identifier in the ECZ database). We use the 2008 list to merge the two datasets with the caveat that the list has not been updated since 2008 to incorporate any new schools that might have been built since then that also qualify as examination centers.⁶

We restrict our data to districts that are relevant for our empirical approach, i.e. where there is a chance of observing a school at the two sides of the rural allowance eligibility threshold. For instance, we drop remote districts where the cutoff is zero and all schools qualify for the allowance (see above), and we drop districts in Copperbelt Province which are very urbanized and where none of the schools are eligible for the allowance.

3.3. Descriptive Statistics

Table 1 presents the descriptive statistics of the administrative data we use. The top panel uses the entire sample of government-run schools from 2004 to 2015 and the bottom panel restricts the sample to the post-treatment period (2011 to 2015) for schools within a 10 Km radius of the

⁶ An additional 708 government-run schools have been added to the Annual School Census between 2008 and 2015. This, however, does not mean that these additional schools are all qualifying examinations centers going by the discussion in the paragraph. It is also worth noting that this does not imply that the 708 schools added to the census are necessarily new schools. They may well be old schools that were never captured in previous census for logistical reasons.

rural hardship allowance threshold, which are the schools for which our effects will be identified. The last 2 rows of the bottom panel show statistics on school eligibility of the allowance and the proportion of schools actually getting the allowance, the latter statistic from the P MEC data described above.

The data in Table 1 shows that, on average, 7% of a school's teaching staff transfer to another school every year. Combining this information with the fact that schools are stocked with about 13 teachers on average (second row) implies that a single teacher leaves every year. The table also shows that the average tenure of teachers at each school is about 10 years, with teachers having 36 years of age and 12 years of schooling on average. The fact that there is little spread in this latter variable around the average is interesting. It does suggest some form of strict enforcement of a rule requiring that teachers must have completed 12 years schooling before they can teach (the formal schooling system in Zambia runs from Grade 1 to Grade 12). Regarding student grades, 14% of boys and 12% of girls sitting for the exams score a Division One. As shown in panel B, schools around the rural allowance eligibility threshold are not very dissimilar from the average, the main difference being that they are somewhat smaller (10 teachers on average as opposed to 13). The last rows panel B show that within a 10kilometer radius of the threshold, 53% of schools qualify to get the allowance whereas 55% actually get the allowance.

To illustrate the rural-urban differences that motivate the introduction of the rural hardship allowance, Table 2 reports differences in teacher and student outcomes across these locations. Notice that the differences we reported are probably a lower bound of the true differences because we have excluded the very rural and very urban districts, as mentioned above. The number of teachers per school tend to be much smaller in rural areas than in urban ones. This might in part be explained by the fact that urban schools are much bigger than rural schools, but is also likely to be driven by hardship concerns, as suggested by the higher percent of transfers out of school and the lower average tenure of teachers shown in the rows below. Moreover, student outcomes are also substantially weaker in rural areas. The only variable

without spatial disparities is the teacher education variable where the average of 12 is equal in both regions - a fact possibly explained by the enforcement of minimum educational requirements to be a teacher.

3.4. Telephone Survey

In order to address the potential problem of teacher-payroll mismatch mentioned above, we conducted a telephone survey where we asked head teachers questions about teachers in their schools. The main questions in the survey asked how many teachers in each school were getting paid from another school and how many teachers were getting the rural hardship allowance. For the survey, we selected from the Annual School Census all schools close to the rural hardship allowance eligibility threshold (within 10kilometers of the threshold) and with a recorded telephone number for the head teacher. To improve power, we stratified the sample into groups of schools close to each other (within 15Kilometers of each other) and drew for each stratum one school at either side of the threshold. This led to a sample of 220 schools and we were able to reach 137 of them, representing 62% of the targeted sample resulting in 88 observations for which we had a full pair. All the head teachers we managed to reach agreed to respond to our queries with the exception of one.

Table 3 shows the descriptive statistics from our sample of the telephone survey. The average number of teachers per school in the telephone survey is somewhat smaller than the corresponding figure from the administrative data in the bottom panel of Table 1. Further, the percentage of schools that qualify for the allowance and the percentage that actually get the allowance in Table 3 are not too different from the ones reported in the bottom panel of Table 1. This suggests that the telephone survey is fairly, though not perfectly, representative. The bias probably comes from the fact that the schools we were able to reach are more likely to have mobile phone coverage, have more economic activity and thus bigger schools.

Table 3 shows that around 40% of teachers are paid from another school, which is line with the findings of payroll mismatch from the Auditor General's Office.⁷ Whereas half of our sample ought to be getting the allowance, 62% state receiving it corresponding to 61% of the teachers. Figures for schools even closer to the threshold (within 5 Km) are similar.

The last two rows of the table show the share of teachers not paid by the school, in schools a priori eligible or ineligible to get the allowance. This can be useful to gauge whether the teacher mismatch is an attempt to "game the system" by teachers that move to an ineligible school. If this were the case, these teachers would seek to remain on the payroll of their previous school to keep the allowance, and we should observe a larger share of teachers not paid by their school in ineligible schools. If, instead, the teacher-payroll mismatch is due to random infrequent updating of the payroll due to lack of capacity, then we should observe little difference between teachers not paid by the school in eligible vs. ineligible schools. The table shows that rather the latter seems to be the case. Both schools eligible and ineligible have similar shares of teachers paid from somewhere else. On the basis of this data, it appears that the teacher mismatch is not simply an attempt to "game the system", at least regarding the rural hardship allowance.

The quality of responses seemed satisfactory. In particular, we asked head teachers two separate questions about the allowance: first, if the school as a whole was eligible to get the allowance and, second, how many teachers, among those paid from the school, were getting the allowance. Theoretically, in an eligible school, all teachers paid from that school ought to be getting the allowance, whereas in a non-eligible school, none should. In a large majority of schools, this was indeed the case, and whenever there were departures, these were small.⁸

⁷ In personal communication with the Auditor General's Office, the percent of teachers getting paid from another school was estimated at 30%.

⁸ In particular, for schools reported to be not eligible, only in one case did the head teacher say that some teachers paid from the school were getting the allowance. For schools reporting to be eligible, only a minority of head teachers provided a different figure for the number of teachers getting the allowance and for those paid by the school, and even then, most were far off by just 1 or 2 teachers.

4. Empirical Approach

Obtaining the pure causal effect of unconditional salaries on teacher outcomes is generally difficult: Schools with highly paid teachers are likely to differ from schools with less well paid teachers in many respects, and all these differences may confound the pure effect of pay. For instance, urban schools may manage to pay higher wages while having better infrastructure, or having students from a wealthier background. In general, it is difficult to distinguish between the role of students' background from the role of teachers' pay on learning outcomes.

The specific way the rural hardship allowance is implemented in Zambia provides us with an opportunity to estimate the effect of wage income on the behavior of teachers purged from any potential bias using a fuzzy regression discontinuity approach. In particular, we can use the eligibility rule based on distance to district center as an instrument for teacher salaries.

Our data is at the school level. Thus, we can estimate the effect of the "wage bill" at the school level on school outcomes, such as the transfer rate, the stock of teachers at a school and the performance of students on the Grade 7 Examinations.

4.1. RDD Model

Consider an outcome Y_{ij} of school i . Each school is observed at a particular time t but we omit the subscript to lighten notation. Denote the per teacher salary income received in school i by W_i . Each district type has a distance cutoff to determine eligibility for the allowance. We denote by d_i the distance of each school to the relevant cutoff. The equation of interest is then

$$Y_i = \alpha + g(d_i) + \rho \log(W_i) + \beta' z_i + u_i \quad (1)$$

where g is a flexible function of the distance of each school to the relevant cutoff and z_i is a vector of control variables. The coefficient of interest is ρ which we assume captures the effect of all the ways in which wages affect outcomes (for instance via attracting teachers with specific characteristics).

Per teacher wages W_i might be correlated with u_i because schools pay an amount of salaries that depend on characteristics linked with the performance of the school. For instance, more ambitious teachers may obtain higher qualifications and thereby obtain higher wages while teaching in areas where students have better family background. For this reason, we exploit variation in the rural hardship allowance across schools. We decompose total wages into the rural hardship allowance and the rest. We denote non-allowance salary by \widetilde{W}_i , the percent increase in income that the allowance implies by r (which is 0.2 in our case) and the share of teachers in the school getting the allowance by n_i . Thus per teacher salary equals $W_i = \widetilde{W}_i(1 + rn_i)$. Applying the approximation $\log(1+x) \approx x$ allows us to rewrite equation 1 as:

$$Y_i = \alpha + g(d_i) + \rho \log \widetilde{W}_i + \rho r n_i + \beta' z_i + u_i \quad (2).$$

This implies that our coefficient of interest ρ can be identified by exogenous variation in the share of teachers that obtain the allowance. Conditional on a smooth function of d_i , eligibility to obtain the allowance ought to be random. This is plausible given that the eligibility rule is based on distance to the nearest district center as computed by Ministry of Lands using a school's and district center GPS coordinates. Therefore, there are scant possibilities of manipulating eligibility to obtain the allowance (we formally test the possibility of manipulations in the running variable below). Our first stage is:

$$n_i = \alpha_1 + h(d_i) + \tau I\{d_i > 0\} + \gamma' z_i + v_i \quad (3)$$

where $I\{d_i > 0\}$ is an indicator of whether the school is beyond the distance cutoff $d=0$. The coefficient τ will be less than one to the extent that the allocation rule is not strictly followed, either because some ineligible schools do get the allowance or because of the payroll mismatch issue.

In our administrative data (specifically from P MEC), we only have information on whether the school as a whole gets the allowance or not. We do not have information on the share of teachers getting the allowance. Therefore, strictly speaking, we can only obtain τ and consequently ρ using the first stage from the telephone survey described in Section 3.⁹ In the event, the first stages from the administrative data and the telephone survey are not very different. Therefore, for convenience, we report our IV results using the administrative data and discuss potential necessary adjustments where appropriate. The interpretation of the IV coefficients that we will report is thus the effect of having all teachers in the school (as opposed to none) obtain the allowance, which is equivalent to the effect of an increase in the teacher wage bill of 20%.

We estimate different specifications of the equations above using alternative windows (bandwidths) around the cutoff with different polynomial specifications of the running variable. We report results for three bandwidth specifications: optimal bandwidth, 10 Km and 20 Km. The optimal bandwidths are chosen using the method described in Calonico et al. (2017a) which, among other things, incorporates and improves upon the well known method of Imbens and Kalyanaraman (2012). For the optimal and 10kilometer bandwidths, we use a linear polynomial of the running variable allowed to differ at either side of the threshold (i.e. interacted with the

⁹ Using administrative data to perform the first stage is still informative. The instrumental variables (IV) estimate using this first stage captures the effect of the policy as intended. Policymakers may use some discretion and grant the allowance to schools that are ineligible by the distance criterion alone. The reduced form result does not take this into account and would therefore underestimate the intended effect of the policy. The IV estimate using the administrative data as first stage takes account of this discretion. However, the teacher payroll mismatch is unintended. If, in the limit, the mismatch was complete, schools at either side of the threshold would have the same share of teachers getting the allowance. The appropriate assessment of the policy would be that it failed, because the policy would indeed have failed to reward teachers in rural areas. In this scenario, the reduced form coefficient and the IV coefficient using administrative data would show a zero effect. But the IV using the telephone survey would be indeterminate because the first-stage coefficient would be zero.

indicator variable $I(d_i > 0)$). For the 20 Km bandwidth, we use a polynomial of degree 3 with no interaction.

We use specifications without additional controls as well as specification that control for district and for the corresponding pre-treatment outcome.

4.2. Manipulation of the running variable

The RDD approach rests on the fact that there is random assignment into treatment and control groups within a neighborhood of the eligibility threshold. In other words, subjects should not manipulate their way into either the treatment or control group. In our particular case, manipulation would entail that schools falsify their GPS coordinates so as to be further from the nearest district center than they really are. The probability of this happening is very small for two reasons. First, GPS coordinates are not supplied by individual schools but transparently collected using handheld GPS devices by officers from the Surveyor General's Department in the Ministry of Lands. The Ministry of Lands is, in as far as this is concerned, totally independent from the Ministry of General Education. Second, GPS coordinates lend themselves to easy verification making falsification highly unlikely. In any case, we formally test the possibility that schools are non-randomly sorting themselves into the treatment group by, among other things, manipulating their GPS coordinates.

Figure 1 is a histogram of schools plotted against the running variable (distance to threshold) restricted to 20kilometers around the threshold. Schools with a positive running variable qualify to receive the rural hardship allowance whereas those with a negative running variable do not qualify. If manipulation were present, we would expect to see an unusually high number of schools immediately to the right of zero in our histogram. The fact that we do not see this unusual piling up is suggestive of the absence of manipulation of the running variable.

Manipulation can formally be tested using statistical methods. We use the manipulation test proposed by Cattaneo *et al.* (2017b) which is related to the well known test by McCrary (2008).

In both cases, the null hypothesis is that there is a continuity in the estimated probability density function at the threshold. The advantage of Cattaneo et al.'s test over McCrary's test is that the latter requires subjective pre-binning of the data whereas the former is fully data-driven in this respect. Second, traditional manipulation tests such as McCrary's treat the estimation of the density function as a two-sample problem where density functions on either side of the threshold are estimated independently. A case can be made that cumulative distribution functions (cdfs), from which density estimators are derived, and higher-order derivatives of the density functions themselves are often equal in practice for both groups at the threshold. Incidentally, assuming equal cdfs and equal higher-order derivatives for density functions improves the power of the test.

The absence of pre-binning and the assumption of equal cdfs and higher-order derivatives of the density function leads us to favour Cattaneo et al.'s test. Performing the test on our running variable delivers a p-value of 0.870 (Wald t-statistic = 0.152) hence a failure to reject the null of a continuity at the threshold.¹⁰

4.3. Pretreatment Balance

In addition to the absence of manipulation, the RDD method requires that there be balance in outcomes between control and treatment groups prior to treatment. In our case, this requires that there should not be any significant treatment effects in all our outcome variables prior to 2010.

Table 4 shows the results of such an exercise. Pre-treatment balance is confirmed for all our outcome variables in all specifications with the exception of column 2 for the stock of teachers. The treatment effect is, however, not robust to different specifications. Any treatment effects ought to be robust to different bandwidth choices and different specifications, which is not the case for the stock of teachers variable prior to treatment.¹¹

¹⁰ The bandwidth is restricted to 20kilometers around the zero threshold.

¹¹Unfortunately, we cannot check for pre-treatment balance for the Grade 7 Examinations because we do not have the data for the years before treatment.

5. Results

5.1. First stage

5.1.1. Administrative data

Table 5 shows the first stage coefficients using administrative data. That is, data from the Payroll Management Establishment Control (PMEC) and from the Annual School Census. The first two columns in each table use an optimal bandwidth while the latter two columns respectively use a 10kilometer and 20kilomter radius around the threshold. In some specifications, we control for district fixed effects, where possible for predetermined outcomes and for higher order polynomials of the running variable.

The table shows that schools that are near the threshold and qualify to get the allowance are around 40% more likely to actually get the allowance. This finding is robust to different bandwidth choices and to the inclusion of district dummies and predetermined variables. The F-statistics for instrument strength are all greater than 10. Figure 2 illustrates this jump in schools getting the allowance once the threshold is crossed. The figure also shows that some schools that are close to the threshold but do not qualify nonetheless get the allowance. This is probably due to the type of discretion discussed above.

5.1.2. Telephone Survey data

Table 6 shows the first stage coefficients using the telephone survey. As discussed above, we consider two endogenous variables: school getting the allowance and the share of teachers getting the allowance in a school. We regress each of these endogenous variables on the full set of stratification dummies and the indicator variable of eligibility $d > 0$, using a bandwidth of 10 Km in the first column and of 5 Km in the second. The specification with 10 Km bandwidth controls for distance to the threshold interacted with the eligibility variable. Schools close to the

threshold but at the far end are around 45% more likely to get the allowance.¹² When considering the share of teachers getting the allowance, the coefficient drops slightly to around 0.4.

Under the RDD assumption that schools at either side of threshold are on average similar except for the allowance, this implies that teachers on the other side side of the cutoff earn around 8% more than teachers at non-qualifying schools (0.4 more teachers getting 0.2 more income from the allowance).

5.2. Reduced form results

Tables 7 and 8 show the results of regressing outcome variables directly on the indicator variable of eligibility $d > 0$ while controlling for different specifications of the running variable, district dummies and, where possible, predetermined outcomes. The results in Table 7 refer to outcomes from the Annual School Census and those in Table 8 are for the Grade 7 examinations.

The tables show that generally there are no statistical significant effects of the allowance on teacher or student outcomes. However, beyond this general pattern, there are some specific patterns that deserve to be noted. First, coefficients for the stock of teachers and teacher tenure are almost always positive and occasionally become statistically significant (although only mildly so). Second, however, coefficients for teacher transfers are very small. Third, coefficients for teacher characteristics and student grades are also generally small.

These results suggest that the allowance does not have a strong impact on teachers or students. The allowance may have some impact in keeping teachers in rural areas, but the evidence in support of this is rather weak. While the evidence on the stock of teachers and teacher tenure

¹² This is slightly higher than the estimates from the administrative data, possibly because the schools in the telephone survey are somewhat more “developed” than average, as mentioned above, and may implement the rule more stringently.

at the school is consistent with this, the small coefficients for transfers is not.¹³ In any case, what does seem fairly clear is that the allowance does not seem to have a significant effect on teacher characteristics or on student performance.

5.3. IV results

Tables 9 and 10 show the IV results obtained from using the indicator variable of eligibility $d > 0$ as an instrumental variable for schools actually getting the allowance. The tables respectively report results for teacher characteristics and Grade 7 examinations.

As expected, IV results are qualitatively the same as the reduced form ones, with almost no statistically significant result. The advantage is that they can be interpreted quantitatively more easily, as the effect of having all teachers obtain the rural allowance (as opposed to none), which is equivalent to a 20% salary increase. Most of the coefficients are small, notably regarding teacher transfers, teacher characteristics and student grades. The coefficients for the stock of teachers and teacher tenure, however, are non-negligible. The coefficients imply that a school obtaining the rural allowance would increase teacher tenure by one year on average and succeed in retaining around one more teacher in stock. Again, however, these estimates are very imprecise, particularly those for the teacher stock, which varies quite substantially from specification to specification. As before, the conclusion that seems to emerge from the table is that the allowance may have some impact on teacher retention, although the evidence for this is weak, but has little effect on other teacher and student outcomes.

6. Discussion

Our null results regarding the effect of the allowance on student grades is consistent with recent findings in the literature on the effect of unconditional salary increases on teacher performance and student learning (Andrade da Silva Filho et al. 2014, Cabrera et al 2016, Das et al. 2016 and

¹³ As mentioned above, the transfer variable is a noisy measure of hardship-related departures from the school because this variable incorporates other motives for transfers and because it does not include resignations out of the profession. Miss-measurement could be a reason why the results for transfers are not in line with the patterns found for teacher stock and tenure.

De Ree et al. 2015). It appears that this results generalizes to a low income setting such as Zambia, and using a quasi-experiment based on a real policy.

It is somewhat surprising, however, that we find no significant effects of the allowance on the number of teachers nor on teacher transfers and tenure. We do find some generally positive effects on the number of teachers and on teacher tenure, but these effects are almost never significant, and moreover we find virtually no effect on teacher transfers to other schools. This is somewhat surprising because this margin is “incentivized” in the sense that the higher wages coming from the allowance should directly increase the opportunity cost of leaving the school.

There are several possibilities why we may find a low effect of the allowance on teacher departures decisions. First, it may be that we do not have enough power to detect actual existing effects. The allowance increases salaries by 20%, which is a non-negligible amount. However, the discretion in implementing the rule in addition to the teacher-salary mismatch reduces the strength of our instrument, leading to a salary increase of on around 8%. As a result, our IV estimates are fairly imprecise.

Second, it could be that the effect of salaries on teaching mobility decisions is conditional on other variables such as infrastructure (electricity, water, etc.), or distance to amenities. It seems plausible that infrastructure, distance to amenities, and salary income are to a certain extent complements in the utility function of teachers. When one of these variables is very low, the marginal impact of the others may be low as well. In other words, the “binding constraints” for keeping teacher’s in rural areas may be infrastructure and distance to amenities rather than salary income.

The telephone survey we conducted provides an opportunity to shed some light on this question. In particular, we can provide some suggestive evidence on the importance of teacher salaries relative to other factors potentially relevant for teacher mobility.¹⁴ We asked head teachers how many teachers had left the school in the previous 3 years and how many had left

¹⁴ Instead of simply looking at the relative importance of different factors one could study the interaction between salaries and these other variables. In our case, this is impractical because of the small sample size of the survey and the relatively low variation in terms of infrastructure and distance to amenities (the schools we surveyed are around the threshold of being considered rural). Such regressions do not yield any interpretable results.

in order to work at a “better school”. We also considered four types of factors that could potentially affect the decision of teachers to stay. First, we asked about access to infrastructure for teachers in terms electricity and piped water in their dwelling and whether their housing was made out of brick. Second we considered distance to amenities, operationalized as distance to the nearest bank. Third, we considered that not only the amount of salaries matter, but also whether they are paid on time or not, so we asked whether there were delays in paying salaries. Finally, we considered community incentives such as land gifts as potentially stronger triggers of reciprocity than salaries.

Table 11 shows the distribution of these variables in our sample of schools, all of which are close to the threshold for the rural hardship allowance. An average of 30% of the current teacher body left the school in the previous 3 years, one third of them to move to a “better school”. Most teachers are reported to live in brick houses, although only 3% have access to piped water. The average distance to the nearest bank is around 30 Km, consistent with our expectations given that these schools are at the threshold of receiving the allowance. Most of schools reported “a little” delay in receiving salaries in the past 3 years, whereas a very small minority reported large delays. And around two thirds of schools reported that land was given to teachers to farm.

We perform straightforward OLS regressions of teacher departure variables on the explanatory variables just described, combining the three infrastructure variables into an index from a principal component analysis. For comparability, we standardized all variables dividing them by their respective standard deviation. To these explanatory variables, we add the share of teachers in the school obtaining the rural allowance, which is our best measure of salary differences between schools. Table 12 shows the results of this illustrative exercise using as outcome variable total teacher departures in column 1 and those leaving for a “better school” in column 2. Generally, coefficients have the expected sign. And indeed, the allowance variable appears negatively related to both teacher departure variables. However, most coefficients are fairly small and statistically insignificant. The exceptions are distance to nearest bank (for total departures) and payment delays (for departing to a better school). This suggests the possibility that salary considerations maybe secondary to distance to amenities or to delays in payments

for these types of school. Of course, this is a mere conjecture at this stage, which would need to be scrutinized in further research.

7. Concluding Remarks

This paper has studied the effect of unconditional wage increases on teacher and student learning outcomes in Zambia. The rural hardship allowance for teachers in Zambia represents a 20% increase and is allocated in a way that allows us to estimate this effect using a regression discontinuity design. The rule is partially followed. Partly this is because there is some idiosyncrasy around granting the allowance to schools above and beyond the allocation rule, and partly because there is some teacher-salary mismatch whereby teachers get paid from schools other than where they actually teach. Our first stage is nevertheless highly significant and represents a jump in salary of around 8%.

We find no significant effect of the allowance on student test results. More surprisingly, we find not statistically significant effects of the allowance on teacher mobility decisions. There is some consistent positive effect of the allowance on the stock of teachers and on teacher tenure but the effects are often far from significant. We discuss several reasons why this may be the case. We provide suggestive evidence that non-monetary considerations such as distance to amenities or delays in payment may be more relevant for teacher mobility in rural schools than monetary income. This evidence, however, is only illustrative and needs to be explored further in future research.

References

- Auditor General, 2014. Report of the Auditor General on the deployment of teachers in Basic Schools
- Bau, Natalie, and Jishnu Das. 2016. "The Misallocation of Pay and Productivity in the Public Sector: Evidence from the Labor Market for Teachers."

Besley, Timothy J., and Maitreesh Ghatak. 2014. "Solving Agency Problems: Intrinsic Motivation, Incentives, and Productivity." *Background Paper for the 2014 World Development Report*.

Cabrera, José María, and Dinand Webbink. n.d. "Do Higher Salaries Yield Better Teachers and Better Student Outcomes?"

Calonico, S., Cattaneo, M.D., Farrell, M.H & Titiunik, R., 2017a. rdrobust: Software for regression-discontinuity design. *Stata Journal*, 17(2): 372 - 404

Calonico, M.D., Jansson, M. & Ma, X., 2017b. rddensity: Manipulation testing based on Density Discontinuity. *Stata Journal*

Chetty, Raj, John N. Friedman, and Jonah E. Rockoff. 2014. "Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood." *American Economic Review* 104 (9):2633–79. <https://doi.org/10.1257/aer.104.9.2633>.

Dal Bó, Ernesto, Frederico Finan, and Martín A. Rossi. 2013. "Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service." *The Quarterly Journal of Economics* 128 (3):1169–1218.

Das, Jishnu, and Natalie Bau. 2014. "The Misallocation of Pay and Productivity in the Public Sector: Evidence from the Labor Market for Teachers." Working Paper.

De Ree, Joppe, Karthik Muralidharan, Menno Pradhan, and Halsey Rogers. 2015. "Double for Nothing? Experimental Evidence on the Impact of an Unconditional Teacher Salary Increase on Student Performance in Indonesia." National Bureau of Economic Research.

Duflo, Esther, and Rema Hanna. 2005. "Monitoring Works: Getting Teachers to Come to School." National Bureau of Economic Research.

Glewwe, Paul W, Eric A Hanushek, Sarah D Humpage, and Renato Ravina. 2011. "School Resources and Educational Outcomes in Developing Countries: A Review of the Literature from 1990 to 2010." NBER Working Papers 17554. National Bureau of Economic Research, Inc.

Government of the Republic of Zambia, 2006. Vision 2030.

Government of the Republic of Zambia, 2011. Sixth National Development Plan: 2011 - 2015.

Henn, S., 2017. The impacts of Weak State Capacity in Rural Africa. Working Paper.

Imbens, G. & Kalyanaraman, K., 2012. Optimal bandwidth choice for the Regression Discontinuity Estimator. *Review of Economic Studies*, 79(3): 933 – 959

Jackson, C Kirabo, Jonah E Rockoff, and Douglas O Staiger. 2014. “Teacher Effects and Teacher-Related Policies.” *Annu. Rev. Econ.* 6 (1):801–825.

Kremer, Michael, Paul Glewwe, and Nauman Ilias. 2010. “Teacher Incentives.” *American Economic Journal: Applied Economics* 2 (3).

McCrary, J., 2008. Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, 142(2): 698 - 714

Mulcahy-Dunn, A., Herstein, J., Dunn, S. & Sevilla, M.A., 2003. “Teacher Attrition and Absenteeism in Zambia”. Report prepared for USAID, Republic of Zambia Ministry of Education & American Institutes for Research.

Muralidharan, Karthik, and Venkatesh Sundararaman. 2011. “Teacher Performance Pay: Experimental Evidence from India.” *Journal of Political Economy* 119 (1):39–77.

Silva Filho, Geraldo Andrade da, and Cristine Campos de Xavier Pinto. 2014. “Higher Salaries, More Teaching, Better Performance?”

Snilstveit, Birte, Jennifer Stevenson, Radhika Menon, Daniel Phillips, Emma Gallagher, Maisie Geleen, Hannah Jobse, Tanja Schmidt, and Emmanuel Jimenez. 2016. “The Impact of Education Programmes on Learning and School Participation in Low-and Middle-Income Countries.”

World Bank, 2014. World Development Indicators Database. Accessed (July 2014).

TABLES AND FIGURES

Table 1: Descriptive Statistics for full and restricted sample

PANEL A					
Variable	Mean	SD	Minimum	Maximum	Observations
Transfer share	0.069	0.146	0.000	1.000	29007
Stock of Teachers	12.997	14.198	1.000	100	29007
Tenure of Teachers (Years)	10.430	4.216	0.000	43.333	25532
Education level of teachers (Years)	11.828	0.494	6.000	19.000	28542
Age of teachers (Years)	36.244	3.429	20.800	63.000	20392
Percentage of Boys with Division One	14.223	16.071	0.000	100.000	8184
Percentage of Girls with Division One	11.858	15.918	0.000	100.000	8174
PANEL B					
Variable	Mean	SD	Minimum	Maximum	Observations
Transfer share	0.095	0.164	0.000	1.000	3488
Stock of Teachers	10.216	6.841	1.000	88.000	3488
Tenure of Teachers (Years)	9.776	2.898	1.000	36.667	3385
Education level of teachers (Years)	11.975	0.190	9.600	13.333	3411
Age of teachers (Years)	36.404	2.919	9.125	52.000	3384
Percentage of Boys with Division One	13.084	15.115	0.000	100.000	2323
Percentage of Girls with Division Two	11.023	15.302	0.000	100.000	2320
School eligible to get allowance	0.531	0.499	0.000	1.000	3488
School gets allowance	0.555	0.497	0.000	1.000	3488

Descriptive statistics from a combination of the Annual School Census, Grade 7 Examinations and Allowance data from PMEC. Panel A refers to the full sample. Panel B refers to a sample restricted to schools within a 10 Km radius of the threshold post-treatment.

Table 2: Descriptive statistics for urban and rural areas

Variable	Urban	Rural	Observations
Transfer share	0.054	0.078	29007
Stock of teachers	23.441	7.928	29007
Tenure of teachers (Years)	11.208	10.04	25532
Education level of teachers (Years)	11.845	11.820	28542
Age of teachers (Years)	36.878	35.905	20392
Percent of Boys with Division 1	16.213	13.283	8184
Percent of Girls with Division 1	14.179	10.760	8174

Descriptive statistics from the Annual Schools Census for urban and rural Zambia.

Table 3: Descriptive Statistics from telephone survey

	10 Km radius		5 Km radius	
	Mean	SD	Mean	SD
Number of teachers	12.64	8.47	12.74	10.54
Share of teachers paid from another school	0.38	0.24	0.41	0.25
School a priori eligible to get the allowance	0.5	0.5	0.5	0.51
School gets allowance	0.62	0.49	0.62	0.49
Share of teachers getting the allowance	0.61	0.38	0.6	0.37
Share of teachers paid from another school in a priori eligible	0.37	0.23	0.44	0.27
Share of teachers paid from another school in a priori not eligible	0.39	0.25	0.37	0.23
N	88		34	

Descriptive statistics from the telephone survey

Table 4: Pre-treatment balance

	(1)	(2)	(3)	(4)
Share of Transfers	0.253 (0.846) [4651]	0.203 (0.792) [4651]	0.152 (0.822) [4409]	-0.080 (0.763) [8934]
Stock of teachers	1.431 (0.863)* [2339]	1.910 (0.883)** [2339]	0.982 (0.655) [4409]	-0.759 (0.709) [8934]
Teacher Tenure	-0.188 (0.466) [3763]	-0.244 (0.418) [3763]	-0.190 (0.441) [3600]	-0.279 (0.389) [7286]
Teacher Education	0.022 (0.040) [5356]	0.013 (0.038) [5356]	0.015 (0.043) [4379]	0.011 (0.039) [8833]
Teacher Age	0.020 (0.451) [2534]	-0.055 (0.407) [2534]	0.127 (0.452) [2192]	-0.058 (0.402) [4437]
Bandwidth	Optimal	Optimal	10	20
Poly. Order	1	1	1	3
District Controls	No	Yes	Yes	Yes
Interaction term	Yes	Yes	Yes	No

Robust standard errors clustered at the school level are reported in parenthesis. Sample sizes are reported in square brackets. Significance codes: 0.01 '***' 0.05 '**' 0.1 '*'. Each row uses a different outcome variable. The first two columns use an optimal window around the threshold based on the method in Calonico *et al.* (2017b). The last two columns use larger windows of respectively 10kms and 20kms around the threshold. The interaction term interacts treatment status with the running variable.

Table 5: First stage administrative data

	(1)	(2)	(3)	(4)
Allowance	0.341 (0.080)*** [2559]	0.310 (0.081)*** [2559]	0.412 (0.065)*** [3508]	0.448 (0.059)*** [7116]
F-Stat	18.170	14.590	40.000	61.49
Bandwidth	Optimal	Optimal	10	20
Poly. Order	1	1	1	3
District Controls	No	Yes	Yes	Yes
Interaction term	Yes	Yes	Yes	No

Robust standard errors clustered at the school level are reported in parenthesis. Sample sizes are reported in square brackets. Significance codes: 0.01 '***' 0.05 '**' 0.1 '*'. Each row uses a different outcome variable. The first two columns use an optimal window around the threshold based on the method in Calonico *et al.* (2017b). The last two columns use larger windows of respectively 10kms and 20kms around the threshold. The interaction term interacts treatment status with the running variable. The F-statistic is associated with the hypothesis that the coefficient is zero.

Table 6: First stage telephone survey

	{1}	{2}
School getting allowance	0.412 (0.205)*	0.529 (0.125)***
Share teachers getting allowance	0.359 (0.132)***	0.421 (0.076)***
Bandwidth	10 Km	5 Km
N	88	34

Robust standard errors clustered at the school level are reported in parenthesis. Significance codes: 0.01 '***' 0.05 '**' 0.1 '*'.

Table 7: Reduced Form, teacher characteristics

	(1)	(2)	(3)	(4)
Share of Transfers	-0.006 (0.015) [2545]	0.003 (0.013) [2533]	0.002 (0.011) [3477]	0.009 (0.011) [7046]
Stock of teachers	1.857 (0.954)* [1754]	0.398 (0.523) [1742]	-0.262 (0.394) [3477]	0.659 (0.483) [7046]
Teacher Tenure	0.385 (0.339) [3213]	0.491 (0.324) [3200]	0.530 (0.312)* [3373]	0.247 (0.289) [6842]
Teacher Education	-0.005 (0.023) [3653]	-0.010 (0.023) [3640]	-0.008 (0.024) [3400]	0.006 (0.022) [6890]
Teacher Age	0.432 (0.400) [2692]	0.445 (0.354) [2678]	0.237 (0.311) [3369]	-0.002 (0.285) [6838]
Bandwidth	Optimal	Optimal	10	20
Poly. Order	1	1	1	3
District Controls	No	Yes	Yes	Yes
Interaction term	Yes	Yes	Yes	No
Predetermined Controls	No	Yes	Yes	Yes

Robust standard errors clustered at the school level are reported in parenthesis. Sample sizes are reported in square brackets. Significance codes: 0.01 '***' 0.05 '**' 0.1 '*'. Each row uses a different outcome variable. The first two columns use an optimal window around the threshold based on the method in Calonico *et al.* (2017b). The last two columns use larger windows of respectively 10kms and 20kms around the threshold. The interaction term interacts treatment status with the running variable.

Table 8: Reduced form Grade 7 examinations

	(1)	(2)	(3)	(4)
Division One- Boys	0.015 (0.018) [1722]	0.004 (0.017) [1722]	0.000 (0.014) [2323]	-0.021 (0.014) [4621]
Division One - Girls	-0.010 (0.017) [1722]	-0.011 (0.018) [1722]	-0.020 (0.015) [2320]	-0.033 (0.014)** [4618]
Bandwidth	Optimal	Optimal	10	20
Poly. Order	1	1	1	3
District Controls	No	Yes	Yes	Yes
Interaction term	Yes	Yes	Yes	No
Predetermined controls	No	No	No	No

Robust standard errors clustered at the school level are reported in parenthesis. Sample sizes are reported in square brackets. Significance codes: 0.01 '***' 0.05 '**' 0.1 '*'. Each row uses a different outcome variable. The first two columns use an optimal window around the threshold based on the method in Calonico *et al.* (2017b). The last two columns use larger windows of respectively 10kms and 20kms around the threshold. The interaction term interacts treatment status with the running variable.

Table 9: IV results for teacher characteristics.

	(1)	(2)	(3)	(4)
Share of Transfers	-0.019 (0.046) [2545]	0.011 (0.045) [2533]	0.005 (0.028) [3477]	0.020 (0.025) [7046]
Stock of teachers	6.163 (3.637)* [1754]	1.620 (2.120) [1742]	-0.644 (0.971) [3477]	1.494 (1.108) [7046]
Teacher Tenure	0.974 (0.895) [3213]	1.308 (0.911) [3200]	1.313 (0.816) [3373]	0.564 (0.669) [6842]
Teacher Education	-0.011 (0.054) [3653]	-0.025 (0.054) [3640]	-0.018 (0.060) [3400]	0.014 (0.049) [6890]
Teacher Age	1.168 (1.144) [2692]	1.289 (1.090) [2678]	0.592 (0.788) [3369]	-0.004 (0.655) [6838]
Bandwidth	Optimal	Optimal	10	20
Poly. Order	1	1	1	3
District Controls	No	Yes	Yes	Yes
Interaction term	Yes	Yes	Yes	No
Predetermined Controls	No	Yes	Yes	Yes

Robust standard errors clustered at the school level are reported in parenthesis. Sample sizes are reported in square brackets. Significance codes: 0.01 '***' 0.05 '**' 0.1 '*'. Each row uses a different outcome variable. The first two columns use an optimal window around the threshold based on the method in Calonico *et al.* (2017b). The last two columns use larger windows of respectively 10kms and 20kms around the threshold. The interaction term interacts treatment status with the running variable.

Table 10: IV results Grade 7 Examinations

	(1)	(2)	(3)	(4)
Division One- Boys	0.045 (0.060) [1713]	0.012 (0.056) [1713]	-0.003 (0.034) [2312]	-0.046 (0.032) [4596]
Division One - Girls	-0.033 (0.055) [1713]	-0.036 (0.056) [1713]	-0.046 (0.034) [2309]	-0.073 (0.032)** [4593]
Bandwidth	Optimal	Optimal	10	20
Poly. Order	1	1	1	3
District Controls	No	Yes	Yes	Yes
Interaction term	Yes	Yes	Yes	No
Predetermined controls	No	No	No	No

Robust standard errors clustered at the school level are reported in parenthesis. Sample sizes are reported in square brackets. Significance codes: 0.01 '***' 0.05 '**' 0.1 '*'. Each row uses a different outcome variable. The first two columns use an optimal window around the threshold based on the method in Calonico *et al.* (2017b). The last two columns use larger windows of respectively 10kms and 20kms around the threshold. The interaction term interacts treatment status with the running variable.

Table 11: Descriptive statistics teacher departure-related variables in telephone survey

Statistic	N	Mean	St. Dev.	Min	Max
Share Teachers that Left	137	0.321	0.244	0.000	1.333
Share Teachers that Left for Better School	137	0.100	0.139	0.000	0.667
Share Teachers with Brick House	137	0.605	0.285	0.000	1.500
Share Teachers with Electricity	137	0.124	0.257	0.000	1.000
Share Teachers with Piped Water	137	0.035	0.139	0.000	1.000
Distance to Bank	137	36.079	15.448	1.000	135.000
Delay in Salary Payment Little	137	0.745	0.438	0	1
Delay in Salary Payment Lot	137	0.022	0.147	0	1
Land Gift	137	0.628	0.485	0	1

Table 12: Regressions of teacher departure on potential explanatory variables

	Share Teachers that Left	Share Teachers that Left for Better School
Share Teachers with Brick House	-0.063 (0.095)	-0.071 (0.093)
Share Teachers with Electricity	-0.016 (0.088)	-0.080 (0.086)
Share Teachers with Piped Water	0.137 (0.095)	0.084 (0.093)
Distance to Bank	0.073 (0.088)	0.211** (0.086)
Delay in Salary Payment Little	0.009 (0.087)	0.086 (0.085)
Delay in Salary Payment Lot	0.083 (0.794)	-0.726 (0.777)
Observations	137	137

Note:

*p<0.1; **p<0.05; ***p<0.01

Figure 1: Histogram of schools against running variable

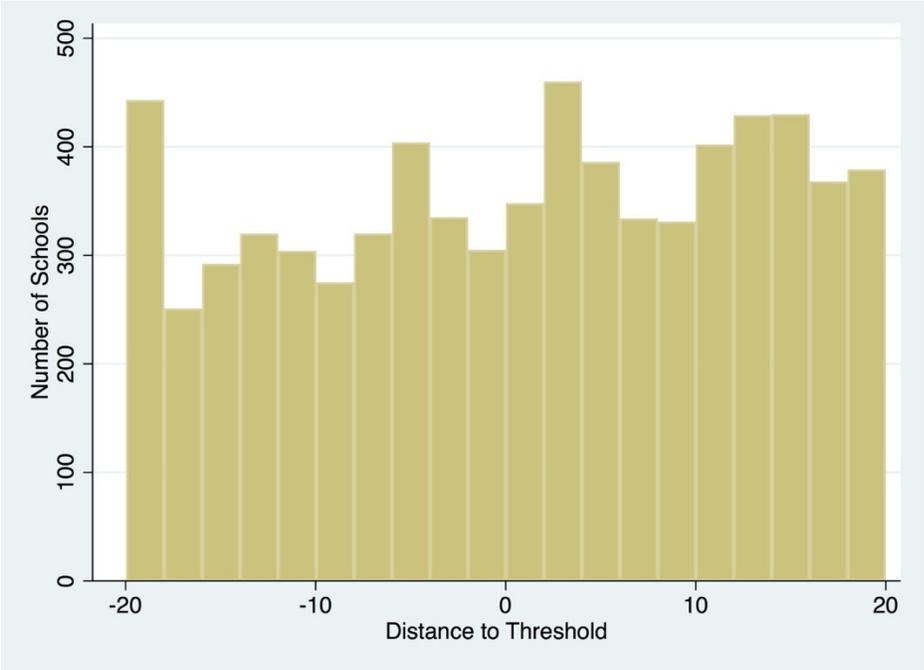


Figure 2: First stage

