

TEACHER PAY AND EDUCATIONAL OUTCOMES: EVIDENCE FROM THE RURAL HARDSHIP ALLOWANCE IN ZAMBIA

GRIEVE CHELWA* , MIQUEL PELLICER†‡ and MASHEKWA MABOSHE*§

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 allowance in Zambia, which corresponds to a salary increase of 20%. This allowance is allocated to schools on the basis of a distance criterion allowing us to use a regression discontinuity design. We use administrative data from 2004 to 2015 on school, teacher characteristics and test scores. The administrative data are complemented with a telephone survey of schools close to the eligibility threshold. We find that crossing the threshold increases the share of teachers obtaining the allowance by 40%. Because of some non-compliance with the allocation rule, our estimates are fairly imprecise. Focusing on provinces with better compliance we find some, albeit weak, evidence that the allowance increases the stock of teachers. We, however, find no effects on teacher characteristics or on student test scores.

JEL Classification: I2

Keywords: Education, policy evaluation, regression discontinuity, Zambia

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 Glewwe *et al.*, 2011; Snilstveit *et al.*, 2016). 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 (Chetty *et al.*, 2018, 2014; Jackson *et al.*, 2014), and recent evidence from Pakistan suggests a high impact of teachers on grades in a developing country setting (Bau and Das, 2017).

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 behaviours. For instance,

* Corresponding author: Senior Lecturer, Graduate School of Business, University of Cape Town, Cape Town, South Africa and Southern African Institute for Policy and Research (SAIPAR), Lusaka, Zambia. E-mail: grievchelwa@gmail.com

† University of Duisburg-Essen, Duisburg, Germany

‡ Southern Africa Labour and Development Research Unit (SALDRU), University of Cape Town

§ Public and Environmental Economics Research Centre (PEERC), University of Johannesburg

teacher presence has been incentivised by monitoring or providing financial incentives conditional on presence (Duflo *et al.*, 2012), and teacher contribution to student performance has been incentivised by providing financial rewards to teachers conditional on positive student test scores (Kremer *et al.*, 2010; Muralidharan and Sundararaman, 2011). 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 (Snilstveit *et al.*, 2016). Moreover, some studies find that such interventions have generated undesirable outcomes such as, for example, teaching-to-the-test (Kremer *et al.*, 2010).

There has been growing interest in the potential role of “soft” approaches to increase teacher productivity that tap into behavioural responses such as reciprocity or intrinsic motivation (Besley and Ghatak, 2014). In particular, higher *unconditional* wages may improve the productivity of public servants via selection of more motivated teachers (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 centres, 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 papers study the effect of unconditional wage increases on teacher and student outcomes in developing countries: in Brazil, Uruguay, Pakistan and Indonesia (Filho and Pinto, 2014; Bau and Das, 2017; Cabrera and Webbink, 2018; de Ree *et al.*, 2018). All these studies tend to find either no effect or a small effect on student performance. However, more evidence is necessary to bring to bear on the question. The studies on Brazil and Pakistan (Filho and Pinto, 2014; Bau and Das, 2017, 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 and Webbink, 2018; de Ree *et al.*, 2018, respectively) address this problem using experimental and quasi-experimental methods. However, the setting of both studies are middle-income countries. Unconditional salary increases may well have different effects in low-income countries where teacher salaries are very low and unconditional salary expansions may be most relevant.

Our paper contributes to this nascent literature using a quasi-experimental approach to study the effect of unconditional salary increments in a low-income setting. Our evaluation of the rural hardship allowance not only concerns a relatively low income country such as Zambia but also schools in rural areas with little infrastructure and amenities. The increase in pay provided by the allowance is 20%, substantially lower than that in de Ree *et al.* (2018). Nevertheless, in a challenging environment such as rural Zambia, an increase of this magnitude can be important.

Our paper also contributes to the literature on improving state capacity in developing countries. And here we concern ourselves with a small literature on whether pecuniary incentives can be effective in attracting and retaining public servants working under difficult conditions. Bo *et al.* (2013) study this issue in Mexico where it is difficult to attract workers to work in remote or challenging municipalities. Exploiting exogenous variation in wages generated by a field experiment, they find that prospective workers are more likely to accept job offers in challenging municipalities if offered higher wages. Antwi

and Phillips (2013) study this within the context of the “Brain Drain” problem in Africa. They take advantage of sudden, policy-induced, wage increases for government health workers in Ghana. They find that the increases dramatically reduce the likelihood that public health workers emigrate from Ghana. This strand of the state capacity literature is surprisingly small given the pervasiveness of under-capacitated states in developing countries, especially in Africa (Mkandawire, 2002), and the role played by remuneration.

Finally, evaluating the rural hardship allowance in Zambia is also important in itself. 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. Our policy evaluation can, therefore, contribute to the design of policies seeking to reduce inequalities in the distribution of educational outcomes in Zambia.

We construct a school-level data set 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 have a list of schools that received the hardship allowance in 2017 provided by the government’s payroll department.

The rule assigning the allowance is based on distances between district centres and schools computed from Global Positioning System (GPS) coordinates. 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 on either side of the threshold are similar.

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 is considered to be a misleading measure of remoteness of the school.¹ This problem reduces our first stage coefficient from 1 (if the rule were implemented perfectly) to around 0.5. Second, there is teacher-payroll mismatch in Zambia (Auditor General, 2014). The payroll database that determines the payment of salaries and allowances sometimes includes teachers at schools where they no longer teach. This implies that there can be teachers obtaining the allowance at schools ineligible to get the allowance. Because there is no official information on the degree of payroll mismatch, we conducted a telephone survey of head teachers for schools around the eligibility threshold stratified in pairs of schools close to each other. We succeeded in obtaining information for 137 schools, corresponding to 44 matched pairs, 1 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 very large in all specifications.

Because of the relatively low levels of compliance with the rule, our estimates are fairly imprecise. For this reason, in follow-up specifications, we conduct the analysis for the most complying provinces, where the first stage coefficient is around 0.65. Overall, we find some, albeit weak, evidence that the allowance increases the stock of teachers.

¹ Schools that *a priori* do not qualify but have a natural barrier (lake, mountain, etc...) in between them and the district center are sometimes granted the allowance.

However, consistent 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.

We also provide suggestive evidence from our telephone survey that distance to amenities and delays in the payment of salaries may be more relevant for explaining teacher attrition around the threshold than the rural allowance.

The paper is organised 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 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 pupil-teacher 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 standardised 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, recognising 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)² have thus focused on improving quality primarily through the large-scale recruitment and retention of teachers. For

² 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.

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 that many teachers might not be retained. 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 (allowances, housing, training, etc.) with the aim of keeping and motivating teachers.

To this end, the government implemented a rural/remote hardship allowance to reduce the attrition of teachers from rural schools. According to the Ministry of Education, in any given year 7% of the teaching staff in rural areas leave versus 3% in urban areas. Similarly, the tenure of teachers in rural schools is on average 2 years shorter than it is in urban schools.³

The allowance first emerged in the 1990s but was of a small quantity and plagued with problems.⁴ 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 than 20 km from the nearest bank, 10 km from the nearest police station, etc.). In 2010, the rule changed and was dramatically simplified. It was decided that the single criterion would be distance to the nearest district centre. 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 between 20 and 25 km and the most urbanised districts had a cut-off of 30 km.

We met several government officials who verified that this rule was actually used to pay the allowance. The allowance is paid directly to teachers by the government's payroll department on the basis 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 centres. Schools are allowed to contest their allocation if, for instance, the school is separated from the nearest district centre by natural barriers (lakes, mountains, etc.) that make 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

³ Attrition and tenure estimates are taken from the Ministry of Education's Annual Schools Census. See data section for more details.

⁴ Personal communication, Ministry of Lands.

is not kept up-to-date. This implies that when teachers move to a new school, they may still appear on the payroll as being part of the former school. This also applies to the rural hardship allowance and implies that eligible schools may have some teachers that do not receive the allowance or vice versa. The extent of the problem is not known, but there are indications that it could be significant. Teacher pay mismatch 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.⁵

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 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.

⁵ Another potential problem for our analysis, raised by a reviewer, is teachers self-selecting into schools at the point of initial deployment. This is a problem because we might erroneously attribute an outcome to the policy when it is the case that teachers with certain characteristics (intrinsically motivated, etc.) are attracted to certain kinds of schools. This is less of a concern in the Zambian context because decisions around where newly qualified teachers are deployed are largely out of the hands of individual teachers. This is because teacher deployment practice in Zambia is largely needs-based. Upon graduating from training colleges and universities, teachers are deployed to schools having deficits in teachers (National Assembly of Zambia, 2017; Zambia Daily Mail, 2017; Zambia Observer, 2018). Given that individual teachers have no influence over which schools are classified as deficit schools, this renders their deployment largely independent of their preferences. This implies that teacher self-selection into schools is less of a concern in the Zambian case.

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 centre ourselves, using the school GPS coordinates and the GPS coordinates of district centres taken from Henn (2016).

The Annual Schools Census has 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.⁶

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 mainly 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 continuously 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

⁶ 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.

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. Scoring a Division One on the exams is equivalent to passing the exams. 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.

To merge the Annual Schools Census with the examinations data, we make use of an attempt in 2008 by 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 centres.⁷

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 either side of the rural allowance eligibility threshold. For instance, we drop remote districts where the cut-off is zero and, therefore, all schools qualify for the allowance. We also drop districts in the Copperbelt Province that are very urbanised and where none of the schools are eligible for the allowance. For our regression discontinuity approach, we focus on schools close to the rural hardship allowance threshold. We keep schools within a 20 km radius of the threshold, which leaves us with a sample size of around 1,500 schools.

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 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 PMEC 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

⁷ An additional 708 government-run schools have been added to the Annual School Census between 2008 and 2015.

Table 1. Descriptive statistics for full and restricted sample

Variable	Mean	SD	Minimum	Maximum	Observations
<i>Panel A</i>					
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
<i>Panel B</i>					
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	20.800	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.

opposed to 13). The last rows in panel B show that within a 10 km 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 motivated 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 per cent 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 for teachers.

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

Table 2. Descriptive statistics for urban and rural areas

Variable	Urban	Rural	Observations
Transfer share	0.054	0.078	29,007
Stock of teachers	23.441	7.928	29,007
Tenure of teachers (Years)	11.208	10.04	25,532
Education level of teachers (Years)	11.845	11.820	28,542
Age of teachers (Years)	36.878	35.905	20,392
Percent of boys with division 1	16.213	13.283	8,184
Percent of girls with division 1	14.179	10.760	8,174

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. km stands for kilometre.

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 10 km of the threshold) and having a telephone number for the head teacher. To improve power, we stratified the sample into groups of schools close to each other (within 15 km 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 one exception.

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 are, therefore, bigger schools.

Table 3 shows that around 40% of teachers are paid from another school, which is in line with the findings of payroll mismatch from the Auditor General's Office (Auditor General, 2014).⁸ 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. If the teacher mismatch originates from teachers moving from rural schools to urban schools without informing payroll, it could threaten the validity of our approach. The table shows, however, that both eligible and ineligible schools have similar shares of teachers paid from somewhere else, suggesting that the problem of teacher mismatch does not affect the validity of our results.

⁸ In personal communication with the Auditor General's Office, the percent of teachers getting paid from another school was estimated at 30%.

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.⁹

4. EMPIRICAL APPROACH

Obtaining the pure causal effect of unconditional salaries on teacher outcomes is generally difficult: Schools with better 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 behaviour 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 centre as an instrument for teacher salaries.

Our data are 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_i 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 cut-off to determine eligibility for the allowance. We denote by d_i the distance of each school to the relevant cut-off. 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 cut-off 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).

⁹ 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.

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 backgrounds. 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 \tilde{W}_i , the per cent 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 $W_i = \tilde{W}_i (1 + rn_i)$. Applying the approximation $\log(1+x) \approx x$ allows us to rewrite equation (1) as:

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

Equation (2) shows that the coefficient of interest ρ , which captures the effect of wages on outcomes in equation (1), 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 centre as computed by Ministry of Lands using a school's and district centre 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. 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 3.¹⁰ In any event, the first stages from the administrative data and the telephone

¹⁰ 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.

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%.¹¹

Our benchmark specification is linear, with a window of 20 km around the cut-off, and allows the slope to differ at either side of the threshold (*i.e.* the running variable is interacted with the indicator variable $I \{d_i > 0\}$). This specification has the potential to yield fairly precise estimates, but could generate bias in our coefficient of interest, since it has a rather large window and does not capture potential non-linearities. To address these concerns, we also report results for two other specifications: a linear specification with a window of 10 km and a cubic specification with a 20 km window. By restricting the window and adding non-linearities, these specifications will feature less bias, but also less precision, and may therefore fail to yield statistically significant results. However, to the extent that the different specifications yield similar estimates, confidence on our results is strengthened. In Tables A2 and A3 in the appendix, we provide robustness checks for the main results using different windows of between 10 and 20 km around the threshold. All our specifications control for constituency fixed effects and, except for Grade 7 results, 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 neighbourhood 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 centre 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 MoGE. 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.

Fig. 1 is a histogram of schools plotted against the running variable (distance to threshold) restricted to 20 km around the threshold. Schools with a positive running

¹¹ Teachers in eligible schools not getting the allowance might believe that, eventually, they will obtain the allowance. This may cause them to stay longer in that school. This would lead to differences in outcomes on the two sides of the threshold even without a difference in having an allowance. If this were the case, the IV would need to be interpreted with caution. We believe, however, that this mechanism, while plausible, is not likely to be substantial. The reason is that there are relatively few eligible schools that do not get the allowance. The biggest reason for non-compliance with the policy is the reverse that schools not eligible are obtaining it. Moreover, from our interviews with key stakeholders, it appears that teachers are generally not aware of the exact boundary of eligibility of the policy and are thus unlikely to respond to the eligibility criterion alone.

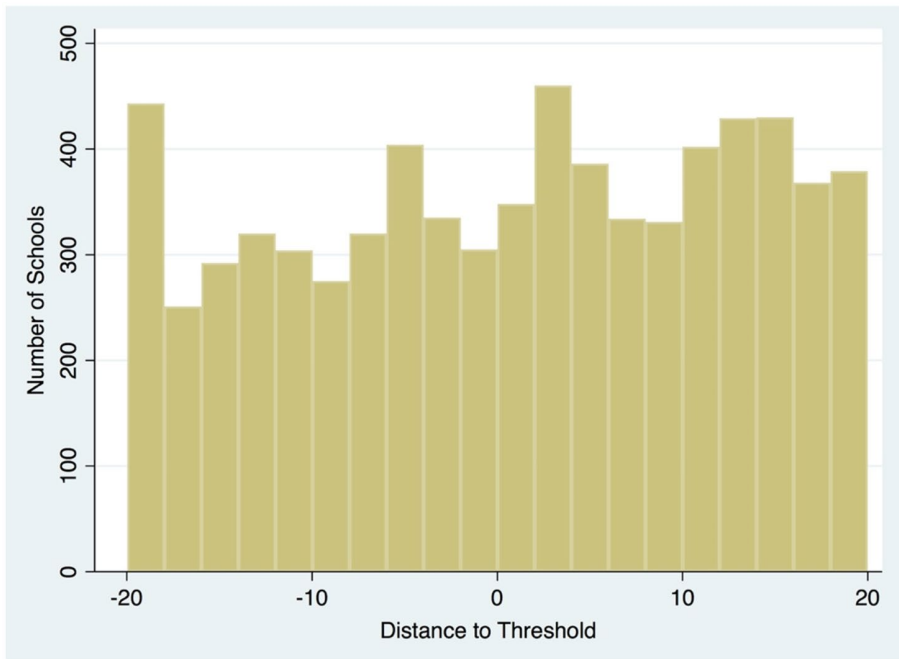


Figure 1. Histogram of schools against running variable [Colour figure can be viewed at wileyonlinelibrary.com]

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 be formally tested using statistical methods. We use the manipulation test proposed by Cattaneo *et al.* (2018) which is related to the well-known test by McCrary (2008). 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.¹²

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. We also check that post-treatment, there is balance in variables not supposed to be affected by the policy, in particular in the number of pupils.

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 log teachers in column 1; *i.e.* the linear specification with a large window.¹³ To assess if this is problematic,

¹² The bandwidth is restricted to 20 km around the zero threshold.

¹³ Unfortunately, we cannot check for pre-treatment balance for the Grade 7 Examinations because we do not have the data for the years before treatment.

Table 4. Pre-treatment balance

	(1)	(2)	(3)
<i>Outcomes</i>			
Log teachers	0.31*** (0.048)	0.044 (0.062)	0.001 (0.058)
Teacher tenure	-0.124 (0.306)	-0.003 (0.457)	-0.134 (0.41)
Share teachers transferred to other school	-0.002 (0.006)	0 (0.009)	-0.002 (0.008)
Teacher age	-0.122 (0.318)	0.176 (0.472)	0.092 (0.426)
Teacher education	0.024 (0.03)	0.013 (0.047)	0.007 (0.042)
Log pupils	0.243*** (0.047)	0.025 (0.061)	0.003 (0.058)
<i>Specifications</i>			
Window	20 km	10 km	20 km
Poly.order	1	1	3
Interaction	Yes	Yes	No
N.schools	1435	698	1435

Robust standard errors clustered at the school level are reported in parenthesis. Significance codes: 0.01 '***' 0.05 '**' 0.1 '*'. Each row uses a different outcome variable. The first two columns use windows of respectively 20 km (km) and 10 km (km) around the threshold and a linear specification, while the third column uses a window of 20 km (km) and a cubic specification. The interaction term interacts treatment status with the running variable.

Fig. 2 plots the graphical counterpart of the regressions for log teachers: the residuals of regressions of log teachers on constituency dummies as a function of the distance to the threshold, using a window of 20 km around the threshold. The left panel shows the predicted values from a regression that uses a linear specification, as in column 1 in Table 4, while the right panel uses a polynomial of degree 3, as in column 3 of Table 4. It is clear that the high and significant coefficient when using a linear specification comes from the non-linearity of the function and not from some genuine difference at either side of the threshold. Because the number of teachers is strongly correlated with the size of the school, it is not surprising to find a similar pattern for (post-treatment) number of pupils: They appear to be unbalanced when considering the benchmark specification, but the differences vanish when considering more flexible specifications. Although this will need to be borne in mind when we consider our results, it does not indicate a basic problem with our approach: Schools at either side of threshold appear to be, as expected, similar on average.

5. RESULTS

5.1 First Stage

(i) *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 column corresponds to the benchmark specification (linear with a 20 km window). The following two columns use more flexible specifications, with a shorter window and a higher order polynomial, respectively.

The table shows that schools that are near the threshold and qualify to get the allowance are around 50% more likely to actually get the allowance with some small

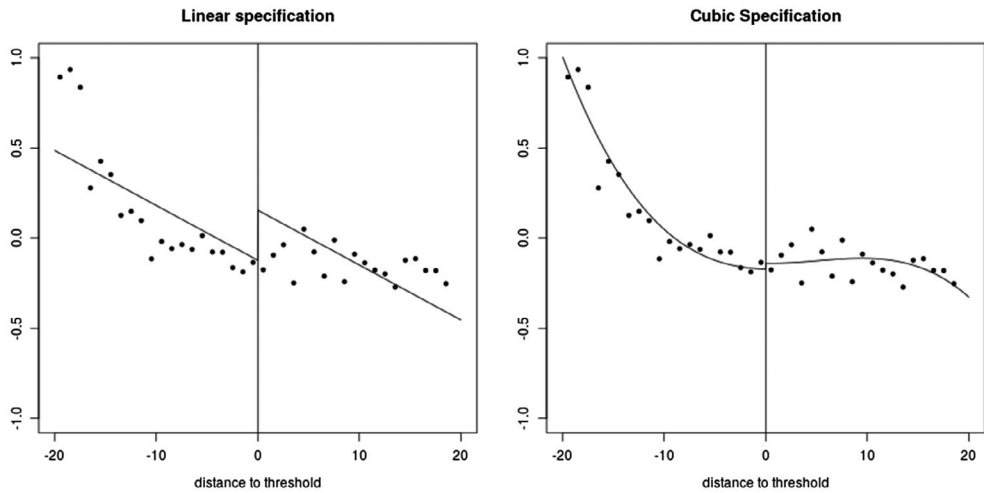


Figure 2. Balance of pre-treatment Log education. Linear and cubic specifications

Table 5. First stage administrative data

	(1)	(2)	(3)
<i>Outcomes</i>			
Allowance	0.55*** (0.042)	0.459*** (0.065)	0.49*** (0.058)
<i>Specifications</i>			
Window	20 Km	10 Km	20 Km
Poly.order	1	1	3
Interaction	Yes	Yes	No
N.schools	1464	714	1464
F statistic	168	50	71

Robust standard errors clustered at the school level are reported in parenthesis. Significance codes: 0.01 '***' 0.05 '**' 0.1 '*'. The first two columns use windows of, respectively, 20 km and 10 km around the threshold and a linear specification, while the third column uses a window of 20 km and a cubic specification. The interaction term interacts treatment status with the running variable.

variations depending on the specification. The F-statistics for instrument strength are all very large, over 50, implying that our instrument is strong.

These first stage coefficients indicate quite a strong degree of non-compliance. To investigate this further, Fig. 3 plots the first stage: the jump in schools getting the allowance once the threshold is crossed. Noncompliance is stronger on the left of the threshold (around 20% non-compliance) than on the right side (around 10%). This makes sense as it implies that some of the schools that do not qualify to get the allowance do so, possibly due to the type of discretion discussed above. The fact that around 10% of schools ought to get the allowance but do not get it is more surprising. Part of this may reflect a partial failure of our merging of payroll data with the Annual School Census data. The telephone survey can help in assessing this and it turns out that the matching between the merged payroll data and the telephone survey data is reasonably good. Only 4% of schools are categorised differently in the payroll data relative to the telephone survey




Figure 3. Illustration of first stage. Share of schools getting the allowance as a function of distance to threshold

data. It appears that there may be other reasons for non-compliance on the right side of the threshold.

In order to investigate the lack of compliance further, we run first stage regressions separately by province. Table 6 shows the results. There are clear differences in the implementation of the rule. Two provinces barely seem to be implementing the rule whereas three others display fairly large coefficients, above 0.6. It thus appears that part of the compliance problem might be genuine problems of implementation of the rule in some parts of the country.

Given the differences in first stage by province, in what follows we perform the analysis for the entire sample, as well as for the provinces with the best compliance.

(ii) Telephone Survey data Table 7 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

Table 6. First stage administrative data, by province

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Allowance	0.355*	0.734***	0.658***	0.63***	0.536***	0.483***	0.469***	0.157
	(0.201)	(0.078)	(0.098)	(0.123)	(0.13)	(0.087)	(0.177)	(0.177)
Subset	Lusaka	Southern	Central	Luapula	Northern/Muchinga	Eastern	Northwestern	Western
N.schools	78	267	216	174	210	316	116	106
Fstat	3	89	45	26	17	31	7	1

Significance codes: 0.01 ‘***’ 0.05 ‘**’ 0.1 ‘*’.

Table 7. First stage telephone survey

	(1)
School getting allowance	0.412 (0.205)*
Share teachers getting allowance	0.359 (0.132)***
Bandwidth	10 km
N	88

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

regress each of these endogenous variables on the full set of stratification dummies and the indicator variable of eligibility $d > 0$. Since we surveyed schools only within a 10 km radius of the eligibility cut-off, we use this bandwidth as benchmark (again, with a linear specification interacted with the eligibility variable). Schools close to the threshold but on the qualifying side are 41% more likely to get the allowance, fairly similar to the estimate using the administrative data with this specification. When considering the share of teachers getting the allowance, the coefficient drops to around 0.36, a fairly small 12% decline.

Under the RDD assumption that schools at either side of the threshold are on average similar except for the allowance, this implies that teachers on the qualifying side of the threshold earn around 7% more than teachers at non-qualifying schools (0.36 more teachers getting 0.2 more income from the allowance). Due to the resulting small sample size, we do not perform the first stage analysis for the telephone survey for the most complying provinces. However, adjusting the 0.65 first stage coefficient for this subsample in the same manner as for the full data, we would obtain a coefficient of around 0.57, which implies a jump in pay of around 11%.

5.2 Intention to Treat Results

Tables 8 shows the results of regressing outcome variables directly on the indicator variable of eligibility $d > 0$ while controlling for different specifications of the running variable, constituency dummies and, where possible, predetermined outcomes.

The table shows that generally there are no statistically 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 log teachers and teacher tenure are almost always positive. Second, however, coefficients for teacher transfers are very small and positive. This is counter to our expectations since if the allowance succeeds in keeping teachers, it should reduce transfers. Third, coefficients for teacher characteristics and student grades are also generally small. For student outcomes, the results are somewhat volatile. Results for boys are generally positive and occasionally statistically significant, but in some specifications it turns negative. For girls, the pattern is exactly the other way around. Given that for student outcomes we are unable to control for predetermined outcomes, these estimates are more imprecise and this is likely to account for the varying results we observe.¹⁴

¹⁴ These results are not essentially altered by considering different bandwidths between 10 and 20 km (see Table A2 in the Appendix).

Table 8. *Intention to treat results, all schools*

	(1)	(2)	(3)
<i>Outcomes</i>			
Log teachers	0.045 (0.026)*	0 (0.036)	0.045 (0.034)
Teacher tenure	0.29 (0.218)	0.405 (0.305)	0.214 (0.289)
Share teachers transferred to other school	0.002 (0.008)	0.004 (0.012)	0.011 (0.011)
Teacher age	-0.054 (0.224)	0.163 (0.316)	-0.01 (0.293)
Teacher education	0.011 (0.016)	-0.003 (0.023)	0.009 (0.022)
Division one boys	0.02 (0.011)*	0.004 (0.015)	-0.014 (0.014)
Division one girls	0.009 (0.011)	-0.019 (0.015)	-0.031 (0.014)**
<i>Specifications</i>			
Window	20 km	10 km	20 km
Poly.order	1	1	3
Interaction	Yes	Yes	No
N.schools	1186	590	1186

Robust standard errors clustered at the school level are reported in parenthesis. Significance codes: 0.01 ‘***’ 0.05 ‘**’ 0.1 ‘*’. Each row uses a different outcome variable. The first two columns use windows of respectively 20 km and 10 km around the threshold and a linear specification, while the third column uses a window of 20 km and a cubic specification. The interaction term interacts treatment status with the running variable.

Thus, 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 at the school is consistent with this, coefficients are generally not statistically significant; moreover, the small positive coefficients for transfers is not consistent with this. In any case, what does seem fairly clear is that the allowance does not seem to have a significant and robust effect on teacher characteristics or on student performance.

5.3 IV Results

Table 9 shows the IV results obtained from using the indicator variable of eligibility $d > 0$ as an instrumental variable for schools actually getting the allowance.

As expected, IV results are qualitatively the same as the intention to treat ones, with almost no statistically significant result. The advantage is that they can be easily interpreted quantitatively, 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 even though they are statistically insignificant. The coefficients imply that a school obtaining the rural allowance would increase teacher tenure by around half a year on average and succeed in retaining 10% more teachers. Again, however, these effects are estimated imprecisely. As before, the conclusion that seems to emerge from the table is that the allowance may not have an effect. However, it is worth cautioning that such a conclusion is arrived at with imprecisely estimated coefficients.

Table 9. IV results, all schools

	(1)	(2)	(3)
<i>Outcomes</i>			
Log teachers	0.099 (0.049)**	0.009 (0.08)	0.094 (0.071)
Teacher tenure	0.507 (0.406)	0.935 (0.711)	0.447 (0.611)
Share teachers transferred to other school	0.003 (0.015)	0.005 (0.027)	0.023 (0.024)
Teacher age	-0.133 (0.419)	0.339 (0.729)	-0.021 (0.619)
Teacher education	0.023 (0.029)	-0.003 (0.054)	0.019 (0.045)
Division one boys	0.037 (0.019)*	0.007 (0.032)	-0.028 (0.028)
Division one girls	0.017 (0.019)	-0.041 (0.032)	-0.062 (0.028)**
<i>Specifications</i>			
Window	20 km	10 km	20 km
Poly.order	1	1	3
Interaction	Yes	Yes	No
N.schools	1186	590	1186

Robust standard errors clustered at the school level are reported in parenthesis. Significance codes: 0.01 ‘***’ 0.05 ‘**’ 0.1 ‘*’. The endogenous variable is an indicator function of whether the school receives the allowance. Each row uses a different outcome variable. The first two columns use windows of respectively 20 km and 10 km around the threshold and a linear specification, while the third column uses a window of 20 km and a cubic specification. The interaction term interacts treatment status with the running variable.

5.4 Results for Most Complying Provinces

Tables 10 and 11 provide the intention to treat and IV results for the three most complying provinces having a first stage coefficient greater than 0.6. Here, the results are slightly more encouraging. Coefficients for log teachers and teacher tenure are now generally larger and, in column 1, statistically significant. The coefficient for teacher transfers is now negative, the “right” sign (though not statistically significant). Results are never statistically significant in the more flexible specifications, with a shorter window and more non-linearity and this makes it impossible to draw strong conclusions. However, some reassurance can be drawn by the fact that even then coefficients remain the “right” sign.¹⁵ Overall, we believe that this paints a suggestive picture that the allowance might be succeeding in either attracting and possibly keeping some teachers in rural areas, at least in the provinces where it is well implemented. Coefficients for teacher characteristics and student performance are, however, not improved when focusing on these provinces and so conclusions regarding these remain unchanged.

It is worth noting that the specification that delivers reasonably statistically significant results (column 1) is the least conservative of all and there is some concern that it might not be flexible enough. Indeed, the balance tests above showed that, for log teachers, this specification fails. However, several reasons suggest that the results in column 1 are not due to lack of flexibility of the specification. First, the coefficient for log teachers is not only positive in column 1, but also in the other specifications where pre-treatment values

¹⁵ Table A3 in the Appendix shows that coefficients remain of the “right” sign even when considering many different bandwidths.

Table 10. *Intention to treat results, most complying schools*

	(1)	(2)	(3)
<i>Outcomes</i>			
Allowance	0.678 (0.057)***	0.546 (0.085)***	0.599 (0.078)***
Log teachers	0.085 (0.043)**	0.033 (0.057)	0.063 (0.054)
Teacher tenure	0.549 (0.308)*	0.308 (0.405)	0.198 (0.4)
Share of transfers	-0.008 (0.011)	-0.006 (0.016)	-0.001 (0.015)
Teacher age	0.17 (0.335)	0.065 (0.486)	0.156 (0.44)
Teacher education	0.004 (0.023)	-0.055 (0.029)*	-0.033 (0.029)
Division one boys	0.023 (0.018)	-0.001 (0.023)	-0.026 (0.022)
Division one girls	0.015 (0.019)	-0.02 (0.024)	-0.043 (0.024)*
<i>Specifications</i>			
Window	20 km	10 km	20 km
Poly.order	1	1	3
Interaction	Yes	Yes	No
N.schools	503	243	503

Robust standard errors clustered at the school level are reported in parenthesis. Significance codes: 0.01 ‘***’ 0.05 ‘**’ 0.1 ‘*’. Each row uses a different outcome variable. The first two columns use windows of respectively 20 km (km) and 10 km (km) around the threshold and a linear specification, while the third column uses a window of 20 km and a cubic specification. The interaction term interacts treatment status with the running variable.

Table 11. *IV results, most complying schools*

	(1)	(2)	(3)
<i>Outcomes</i>			
Log teachers	0.128 (0.064)**	0.056 (0.106)	0.106 (0.093)
Teacher tenure	0.827 (0.474)*	0.572 (0.77)	0.337 (0.685)
Share teachers transferred to other school	-0.012 (0.016)	-0.011 (0.029)	-0.002 (0.025)
Teacher age	0.257 (0.504)	0.155 (0.9)	0.265 (0.751)
Teacher education	0.005 (0.034)	-0.1 (0.056)*	-0.056 (0.049)
Division one boys	0.033 (0.026)	-0.002 (0.041)	-0.045 (0.039)
Division one girls	0.022 (0.027)	-0.035 (0.044)	-0.073 (0.043)*
<i>Specifications</i>			
Window	20 km	10 km	20 km
Poly.order	1	1	3
Interaction	Yes	Yes	No
N.schools	503	243	503

Robust standard errors clustered at the school level are reported in parenthesis. Significance codes: 0.01 ‘***’ 0.05 ‘**’ 0.1 ‘*’. The endogenous variable is an indicator function of whether the school receives the allowance. Each row uses a different outcome variable. The first two columns use windows of respectively 20 km and 10 km around the threshold and a linear specification, while the third column uses a window of 20 km and a cubic specification. The interaction term interacts treatment status with the running variable.

are balanced. Further, the size of the coefficient in column 3 is as large as in column 1. Third, the results in Table 10 already control for pre-treatment log teachers, so the “jump” in pre-treatment log teachers is not driving the results in column 1. Fourth, the balance problem of the specification in column 1 applies only to log teachers. It does not apply to teacher tenure, so its statistical significance (at the 10% level) in column 1 does not suffer from such concern.

Figs 4 and 5 show the graphical representation of the intention to treat results, for log teachers and teacher tenure, respectively. Both figures show the results separately using the entire sample and the data for the most complying provinces. The figures support the points above. There is a jump in log teachers and teacher tenure at the threshold, but the noise in the data obscures it to a certain extent preventing us from stating strong conclusions. At the same time, the jumps do not appear to be driven by non-linearities in the pattern of points, as was the case for pre-treatment log teachers. Therefore, the caution in our conclusions comes from a lack of precision, not from a concern with validity.

The IV results using the most complying provinces do not show a dramatic increase in coefficients relative to the results using the whole country. The allowance is estimated to increase the teacher stock by slightly over 10% and increase the tenure of teachers at the school between 0.5 and 0.8 of a year. Again, however, coefficients are statistically significant only in the benchmark specification.

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 (Filho and Pinto, 2014; Bau and Das, 2017; Cabrera and Webbink, 2018; de Ree *et al.*, 2018). It appears that this results generalises to a low-income setting such as Zambia, using a quasi-experiment based on a real policy.

We find some very suggestive evidence of a positive effect on the number of teachers and on teacher tenure. These are generally not significant when considering the whole




Figure 4. Log teachers, All data and most complying provinces




Figure 5. Teacher Tenure, All data and most complying provinces

country, but become so in some specifications when restricting our attention to provinces where the rule is better implemented.

Our telephone survey, however, provides some suggestive evidence that factors other than teacher salaries may have a stronger impact on teacher mobility decisions. We asked head teachers how many teachers had left the school in the previous 3 years and how many had left 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; we combined these into an index using the first component of a principal component analysis. Second, we considered distance to amenities, operationalised 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. To this end, 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 A1 shows the descriptive statistics of these variables.

We performed a straightforward OLS regression of teacher departure variables on all the explanatory variables just described. For comparability, we standardised all variables dividing them by their respective standard deviation. To these explanatory variables, we added 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. Almost all coefficients are statistically insignificant, which is perhaps not surprising given the small sample size. Even then, the sign and size of the different coefficients provide interesting suggestive evidence. Generally, coefficients have the expected sign. Teacher departures are associated with worse infrastructure, more distance to the nearest bank, delays in payments and, indeed, on being *less* likely to receive the allowance. Interestingly, the largest coefficients are for 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

Table 12. Regressions of teacher departure on potential explanatory variables

	Share teachers that left	Share teachers that left for better school
Share of teachers getting the allowance	-0.063 (0.095)	-0.071 (0.093)
Infrastructure index	-0.016 (0.088)	-0.080 (0.086)
Distance to bank	0.137 (0.095)	0.084 (0.093)
Delay in salary payment	0.073 (0.088)	0.211** (0.086)
Land gift	0.009 (0.087)	0.086 (0.085)

Significance codes: 0.01 ‘***’ 0.05 ‘**’ 0.1 ‘*’.

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 scrutinised 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. This is partly because there is some discretion around granting the allowance to schools 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 7% (11% when considering the most complying provinces).

We find some very weak evidence that the allowance might be achieving its objective regarding teacher mobility decisions: The imprecision of our estimates implies that this conclusion is very tentative. Moreover, we find no significant effect of the allowance on teacher characteristics or on student test results. We provide suggestive evidence that non-monetary considerations such as delays in payment may be more relevant for teacher mobility in rural schools than monetary income itself. This evidence, however, is only illustrative and needs to be explored further in future research.

ACKNOWLEDGEMENTS

We acknowledge funding from the International Growth Center (IGC) at the London School of Economics and Economic Research Southern Africa (ERSA). We would like to thank the following: 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 General Education for availing the Annual School Census data to us, Michael Chilala of the Examinations Council of Zambia for granting us access to the Grade 7 Examinations data, Grevazio Kapanda for able research assistance, Christopher Simusokwe for his help in working with Microsoft Access Files and Aaron Mwewa for helping us understand the teacher landscape in Zambia. We would like to thank seminar/conference participants at the University of Cape Town, Oxford University, University College Cork and SAIPAR in Lusaka for useful comments and suggestions.

REFERENCES

- ANTWI, J. and PHILLIPS, D. C. (2013). Wages and health worker retention: Evidence from public sector wage reforms in Ghana. *Journal of Development Economics*, 102(2013): 101-115.
- AUDITOR GENERAL. (2014). Report of the Auditor General on the deployment of teachers in Basic Schools.
- BAU, N. and DAS, J. (2017). *The Misallocation of Pay and Productivity in the Public Sector: Evidence from the Labor Market for Teachers*. World Bank Policy Research Working Paper No. 8050.
- BESLEY, T. J. and GHATAK, M. (2014). *Solving Agency Problems: Intrinsic Motivation, Incentives, and Productivity*. Background Paper for the 2014 World Development Report.
- BO, D. B., FINAN, F. and ROSSI, M. (2013). Strengthening State Capabilities: The role of financial incentives in the call to Public Service. *Quarterly Journal of Economics*, 128(3): 1169-1218.
- CABRERA, J. M. and WEBBINK, D. (2018). *Do Higher Salaries Yield Better Teachers and Better Student Outcomes?* MPRA Working Paper No. 86972.
- CATTANEO, M. D., JANSSON, M. and MA, X. (2018). Manipulation testing based on density discontinuity. *The Stata Journal*, 18(1): 234-261.
- CHETTY, R., FRIEDMAN, J. N. and ROCKOFF, J. E. (2014). Measuring the impacts of teachers II: Teacher value-added and student outcomes in adulthood. *American Economic Review*, 104(9): 2633-2679.
- DUFLO, E., HANNA, R. and RYAN, S. P. (2012). Incentives work: Getting teachers to come to school. *American Economic Review*, 102(4): 1241-1278.
- FILHO, S. and PINTO, C. (2014). *Higher Salaries, More Teaching, Better Performance?* Sao Paulo School of Economics Working Paper.
- GLEWWE, P. W., HANUSHEK, E. A., HUMPAGE, S. D. and RAVINA, R. (2011). *School Resources and Educational Outcomes in Developing Countries: A Review of the Literature from 1990 to 2010*. NBER Working Paper No. 17554.
- GOVERNMENT OF THE REPUBLIC OF ZAMBIA. (2006). Vision 2030.
- _____. (2011). Sixth National Development Plan: 2011-2015.
- HENN, S. (2016). *The impacts of Weak State Capacity in Rural Africa*. Harvard University Working Paper.
- JACKSON, C. K., ROCKOFF, J. E. and STAIGER, D. O. (2014). Teacher effects and teacher-related policies. *Annual Review of Economics*, 6(1): 801-825.
- KREMER, M., GLEWWE, P. and ILIAS, N. (2010). Teacher Incentives. *American Economic Journal: Applied Economics*, 2(3): 205-227.
- MCCRARY, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, 142(2): 698-714.
- MKANDAWIRE, T. (2002). Incentives, Governance, and Capacity Development in Africa. *African Issues*, 30(1): 15-20.
- MULCAHY-DUNN, A., HERSTEIN, J., DUNN, S. and 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, K. and SUNDARARAMAN, V. (2011). Teacher performance pay: Experimental evidence from India. *Journal of Political Economy*, 119(1): 39-77.
- NATIONAL ASSEMBLY OF ZAMBIA. (2017). Ministerial Statement by the Honourable Minister of General Education. Available at: http://www.parliament.gov.zm/node/6417?fbclid=IwAR1KxzamqH9PKVnunU_oLJfYpAPUIlykAILaAY08mD1x1A0LwYGIiDAUqsc [Accessed April 2019].
- DE REE, J., MURALIDHARAN, K., PRADHAN, M. and ROGERS, H. (2018). Double for nothing? Experimental evidence on an unconditional teacher salary increase in Indonesia. *Quarterly Journal of Economics*, 133(2): 993-1039.
- SNILSTVEIT, B., STEVENSON, J., MENON, R., PHILLIPS, D., GALLAGHER, E., GELEEN, M., JOBSE, H., SCHMIDT, T. and JIMENEZ, E. (2016). The impact of education programmes on learning and school participation in low-and middle-income countries: A systematic review and report. International Initiative for Impact Evaluation (3ie). London.
- WORLD BANK. (2014). World Development Indicators Database. Available at: <https://databank.worldbank.org/data/source/world-development-indicators> [Accessed July 2014].
- ZAMBIA DAILY MAIL. (2017). No excuse for teachers, Opinion Editorial. Available at: <http://www.daily-mail.co.zm/no-excuse-for-teachers> [Accessed April 2019].
- ZAMBIAN OBSERVER. (2018). Ministry of General Education – Teacher Recruitment 2018. Available at: <https://www.zambianobserver.com/ministry-of-general-education-teacher-recruitment-2018> [Accessed April 2019].

APPENDIX

Table A1. 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 A2. Intention to Treat Results, full sample with different bandwidths

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
<i>Outcomes</i>											
Log teachers	0.045* (0.026)	0.05* (0.026)	0.057** (0.027)	0.048* (0.028)	0.049* (0.028)	0.046 (0.029)	0.04 (0.03)	0.04 (0.031)	0.025 (0.032)	0.013 (0.034)	0 (0.036)
Teacher tenure	0.29 (0.218)	0.23 (0.222)	0.193 (0.228)	0.272 (0.234)	0.285 (0.243)	0.309 (0.251)	0.282 (0.262)	0.265 (0.274)	0.282 (0.275)	0.295 (0.29)	0.405 (0.305)
Share teachers transferred	0.002 (0.008)	0.004 (0.008)	0.004 (0.009)	0.007 (0.009)	0.007 (0.009)	0.009 (0.009)	0.008 (0.01)	0.006 (0.01)	0.006 (0.011)	0.003 (0.011)	0.004 (0.012)
Teacher age	-0.054 (0.224)	-0.102 (0.23)	-0.089 (0.237)	-0.09 (0.245)	-0.063 (0.25)	-0.015 (0.255)	-0.09 (0.264)	-0.052 (0.271)	0.065 (0.286)	0.11 (0.299)	0.163 (0.316)
Teacher education	0.011 (0.016)	0.011 (0.016)	0.019 (0.017)	0.006 (0.017)	0.005 (0.017)	0.008 (0.018)	0.012 (0.019)	0.011 (0.02)	0.008 (0.021)	-0.002 (0.023)	-0.003 (0.023)
Division one boys	0.02* (0.011)	0.01 (0.011)	0.001 (0.012)	0.001 (0.012)	-0.001 (0.012)	0.001 (0.012)	-0.001 (0.013)	0 (0.013)	-0.007 (0.014)	-0.002 (0.014)	0.004 (0.015)
Division one girls	0.009 (0.011)	-0.002 (0.011)	-0.012 (0.011)	-0.014 (0.012)	-0.014 (0.012)	-0.011 (0.012)	-0.014 (0.013)	-0.013 (0.013)	-0.022 (0.014)	-0.02 (0.014)	-0.019 (0.015)
<i>Specifications</i>											
W>window	20km	19km	18 km	17 km	16 km	15 km	14 km	13 km	12 km	11 km	10 km
Poly,order	1	1	1	1	1	1	1	1	1	1	1
Interaction	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N.schools	1186	1127	1060	1004	955	891	825	768	702	646	590

Robust standard errors clustered at the school level are reported in parenthesis. Significance codes: 0.01 ‘***’, 0.05 ‘**’, 0.1 ‘*’. Each row uses a different outcome variable. Each column uses a different window around the threshold running from 20 km to 10 km. All columns employ a polynomial of order 1. The interaction term interacts treatment status with the running variable.

Table A3. Intention to Treat Results, most complying provinces with different bandwidths

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Outcomes											
Log teachers	0.085** (0.043)	0.091** (0.041)	0.098** (0.043)	0.074* (0.044)	0.068 (0.044)	0.07 (0.046)	0.069 (0.048)	0.059 (0.049)	0.053 (0.051)	0.036 (0.053)	0.033 (0.057)
Teacher tenure	0.549* (0.308)	0.494 (0.309)	0.444 (0.316)	0.408 (0.325)	0.403 (0.331)	0.322 (0.343)	0.262 (0.353)	0.092 (0.359)	0.146 (0.373)	0.184 (0.394)	0.308 (0.405)
Share of transfers	-0.008 (0.011)	-0.008 (0.011)	-0.007 (0.012)	-0.005 (0.012)	-0.001 (0.012)	0 (0.012)	-0.002 (0.013)	-0.002 (0.013)	0.003 (0.014)	-0.006 (0.015)	-0.006 (0.016)
Teacher age	0.17 (0.335)	0.142 (0.343)	0.159 (0.355)	0.138 (0.373)	0.207 (0.373)	0.244 (0.384)	0.165 (0.399)	0.056 (0.413)	-0.074 (0.436)	0.07 (0.461)	0.065 (0.486)
Teacher education	0.004 (0.023)	0.004 (0.024)	0.004 (0.026)	-0.02 (0.022)	-0.02 (0.023)	-0.018 (0.024)	-0.022 (0.025)	-0.025 (0.026)	-0.033 (0.026)	-0.05* (0.028)	-0.055* (0.029)
Division one boys	0.023 (0.018)	0.01 (0.018)	-0.004 (0.019)	-0.013 (0.018)	-0.013 (0.019)	-0.008 (0.019)	-0.01 (0.02)	-0.013 (0.021)	-0.014 (0.023)	-0.014 (0.022)	-0.001 (0.023)
Division one girls	0.015 (0.019)	0.003 (0.02)	-0.013 (0.02)	-0.025 (0.02)	-0.023 (0.02)	-0.022 (0.021)	-0.024 (0.022)	-0.022 (0.022)	-0.025 (0.024)	-0.026 (0.024)	-0.02 (0.024)
Specifications											
Window	20 km	19 km	18 km	17 km	16 km	15 km	14 km	13 km	12 km	11 km	10 km
Poly.order	1	1	1	1	1	1	1	1	1	1	1
Interaction	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N.schools	503	481	452	426	409	373	349	327	294	268	243

Robust standard errors clustered at the school level are reported in parenthesis. Significance codes: 0.01 ****, 0.05 ***, 0.1 **, Each row uses a different outcome variable. Each column uses a different window around the threshold running from 20 km to 10 km. All columns employ a polynomial of order 1. The interaction term interacts treatment status with the running variable.

Copyright of South African Journal of Economics is the property of Wiley-Blackwell and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.