



Open Access Articles

Incentives for Teacher Relocation: Evidence from the Gambian Hardship Allowance

The Faculty of Oregon State University has made this article openly available.
Please share how this access benefits you. Your story matters.

Citation	Pugatch, T., & Schroeder, E. (2014). Incentives for Teacher Relocation: Evidence from the Gambian Hardship Allowance. <i>Economics of Education Review</i> , 41, 120-136. doi:10.1016/j.econedurev.2014.04.003
DOI	10.1016/j.econedurev.2014.04.003
Publisher	Elsevier
Version	Accepted Manuscript
Terms of Use	http://cdss.library.oregonstate.edu/sa-termsofuse

Incentives for Teacher Relocation: Evidence from the Gambian Hardship Allowance

Todd Pugatch
Oregon State University and IZA

Elizabeth Schroeder
Oregon State University

April 2014

Abstract

We evaluate the impact of the Gambian hardship allowance, which provides a salary premium of 30-40% to primary school teachers in remote locations, on the distribution and characteristics of teachers across schools. A geographic discontinuity in the policy's implementation and the presence of common pre-treatment trends between hardship and non-hardship schools provide sources of identifying variation. We find that the hardship allowance increased the share of qualified (certified) teachers by 10 percentage points. The policy also reduced the pupil-qualified teacher ratio by 27, or 61% of the mean, in recipient schools close to the distance threshold. Further analysis suggests that these gains were not merely the result of teachers switching from non-hardship to hardship schools. With similar policies in place in more than two dozen other developing countries, our study provides an important piece of evidence on their effectiveness.

Keywords: teacher labor markets, rural schools, Gambia, program evaluation, regression discontinuity

JEL codes: I25, I28, J38, J45, J61, O12, O15

1. Introduction

Incentives designed to recruit teachers to rural areas are growing in popularity around the developing world. These policies, intended to improve the quality of education in the most disadvantaged areas, remain largely empirically untested. In 2005, the Gambia introduced a policy to attract qualified (certified) teachers to schools in rural areas. The program provides a salary premium, known as a hardship allowance, to primary school teachers who work in the poorest and most remotely located regions of the country. The premium is large: 30%, 35%, or 40%, depending on distance from the capital. Two features of the policy and its environment allow us to evaluate the impact of the policy on the distribution and characteristics of teachers across schools. First, schools are classified as being in a hardship area if they are located more than 3 kilometers from a main road, generating a geographic discontinuity in receipt of the allowance between otherwise similar schools. Second, schools that did and did not receive the allowance experienced common outcome trends in periods before implementation of the policy, allowing for a comparison of differences in these outcomes after the policy was implemented.

We use these exogenous sources of variation to estimate the effect of the hardship allowance on the presence of qualified teachers in Gambian primary schools. Using difference-in-differences and regression discontinuity specifications, we find that the hardship allowance increased the share of qualified teachers by 10 percentage points. The policy also reduced the pupil-qualified teacher ratio by 27, or 61% of the mean, in recipient schools close to the distance threshold. We conduct the analysis using administrative data on the universe of eligible schools. We verify our identification strategies with a series of tests, and our impact estimates survive a battery of falsification exercises. A back-of-the-envelope calculation suggests that these gains were not merely the result of teachers switching from non-hardship to hardship schools. To our knowledge, this is the first credibly identified evaluation of a geographic recruitment bonus on the presence of qualified teachers for the developing world.¹ With similar policies in place in more than two dozen other developing countries, our study provides an important first piece of evidence on their effectiveness.

¹ An important related study is Urquiola and Vegas (2005), which uses plausibly exogenous variation in the salary premium for rural teachers in Bolivia to evaluate its effect on teacher labor supply. However, they do not distinguish teachers by qualifications, nor do they analyze whether the salary premium attracted additional teachers to rural schools.

The hardship allowance policy reflects concern over the wide discrepancies in educational opportunities between rural and urban areas in developing countries such as the Gambia.² McEwan (1999) documents the many ways in which rural schools are worse off than their urban counterparts in developing areas – lower enrollment and test scores, higher instances of grade repetition and dropouts, as well as shortages of inputs such as skilled teachers. The lack of teacher quality is often advanced as a potential cause of lower performance in rural schools. Evidence that relates observable markers of teacher quality to student performance in developing countries is mixed, however. Glewwe et al. (2011) survey the literature and find that teacher education, in-service training, and salary tend to exert a positive influence on student learning. On the other hand, surveys by Galiani and Perez-Truglia (2011) and McEwan (2013) find that students of contract (unqualified) teachers tend to perform better than their peers.³ This latter evidence, which comes from randomized experiments (Banerjee et al. 2007; Duflo, Dupas, and Kremer 2011; Muralidharan and Sundararaman 2013), runs counter to the conventional wisdom among Gambian policymakers. Regardless of the true effect of teacher certification on student performance, gaining a better understanding of teacher location choices can help disentangle the role of teacher characteristics from other factors that influence outcomes.

Disparities in educational outcomes have led to a growing focus on teacher labor markets in both developed and developing countries. In developed countries, there is evidence that teachers are responsive to wages (see Boyd et al. 2013 for a summary) and school conditions (Barbieri, Rossetti, and Sestito 2011), but are willing to trade off salary for more desirable job characteristics. The literature on compensating differentials attempts to estimate the increase in wages or other benefits that would be required to attract teachers to schools with less desired characteristics, typically using hedonic models. This literature began with Antos and Rosen (1975), who found that white teachers in the United States in 1965 would require \$300 in additional salary to accept the conditions of the average black teacher. The literature continues to find tradeoffs between salary and characteristics such as an urban setting and low-performing students. Hanushek, Kain, and Rivkin (2004) find that teacher mobility responds strongly to race and student achievement in public schools in Texas, and provide estimates of the salary increases that would be required to offset such characteristics as lower student achievement or, for

² Disparities among schools have also been a longstanding concern of developed country policymakers, of course. See Benabou, Kramarz, and Prost (2009) for an evaluation of a policy to direct additional resources to disadvantaged schools in France, for instance.

³ These studies are consistent with evidence from the United States that teacher qualifications are not associated with gains in student performance (Darling-Hammond, Berry, and Thoreson 2001; Hanushek 2003; Rivkin, Hanushek, and Kain 2005).

white teachers, higher proportions of minority students. Recently, Boyd et al. (2013) use a two-sided matching model, and find that elementary school teachers in New York State have preferences for schools that are geographically close to them, are located in suburban areas, and have lower percentages of poor students.

Such findings have led to interest in experimenting with policies to attract teachers to urban areas with wage premiums. There is very little evidence on the effectiveness of these policies, however. Clotfelter et al. (2008) use hazard models to evaluate the impact of a program in North Carolina on teacher retention. Annual bonuses were awarded to teachers who served in low-income or low-performing schools, and the authors find that the bonus reduced turnover rates by 17%. Similarly, Hough (2012) uses difference-in-differences to find that a salary increase in an urban district in San Francisco improved the size and quality of the teacher applicant pool.

The idea of offering wage premiums or other benefits to recruit teachers to underserved areas has become perhaps even more popular in developing countries. Mulkeen and Chen (2008) document the types of challenges facing teachers who live and work in rural areas in a selection of Sub-Saharan African countries. These include harsh living conditions, professional isolation, and lack of interest in education within the community, as well as differences in local languages and ethnicities. Female teachers are especially unwilling to work in rural areas, since single women may feel unsafe in an isolated rural area, while married women may be unable to bring their families and spouses due to fewer job opportunities for family members.

We found documentation of over 40 policies designed to recruit teachers to rural areas in 29 different developing countries.⁴ Of these, 16 were pecuniary awards, making some type of hardship allowance the most popular policy. Despite their growing popularity, however, there is little to no empirical evidence of whether such policies are effective at recruiting teachers to underserved areas. A variety of barriers might diminish the impact of incentives for teachers to move, perhaps the most important of which is limited mobility. Evidence shows that teacher labor markets are geographically small, both in developed countries (Boyd et al. 2005) and developing countries (Jaramillo 2012). Jaramillo finds that having been born in a given province in Peru is a strong determinant of taking a first teaching job in that province. These results highlight the fact that a hardship allowance might work through two different channels. If the policy does not encourage existing qualified teachers to move to

⁴ A list of these policies is omitted for brevity but available upon request.

rural areas, it might attract more individuals, especially those born or living in rural areas, into the teaching profession.

Given the quasi-experimental variation in its implementation, the Gambian hardship allowance provides a nearly ideal opportunity to test the effectiveness of a hardship allowance. Moreover, the ability to conduct the analysis using two separate methods (difference-in-differences and regression discontinuity) offers several advantages. First, the methods estimate different parameters: difference-in-differences estimates the effect of the policy in all hardship schools, while regression discontinuity limits attention to schools just beyond the distance threshold. We expand upon these issues later in the paper. Second, the methods estimate effects over different time horizons. Difference-in-differences averages the policy's effects in the pre- and post-treatment periods, while regression discontinuity focuses on long-run effects because we use data collected 8 years after the policy was introduced. Finally, using two distinct methods provides a built-in check for mutual consistency of results.

The paper proceeds as follows. Below in Section 2, we describe the program and its administration. Section 3 discusses our identification strategy and methodology. Section 4 describes the data and Section 5 presents results. Section 6 concludes.

2. Program Description⁵

Primary and secondary schools in the Gambia are managed by the Ministry of Basic and Secondary Education (MoBSE). MoBSE divides the country into six administrative regions. Schools are classified as Lower Basic (LBS, grades 1-6), Basic Cycle (BCS, grades 1-9), Upper Basic (UBS, grades 7-9), or Senior Secondary (SSS, grades 10-12). School administrations are classified as government, grant-aided (private but funded or subsidized by government), or private. Teachers at these schools are classified as “qualified” if they have completed the three-year teacher training course offered by Gambia College, and “unqualified” otherwise. Qualified teachers receive substantially higher salaries than unqualified teachers (2.5 times at entry, with further seniority increases possible [Mulkeen 2007]) and are more sought after by school administrators for placements.

Each administrative region allocates teachers among its schools based on a pool of candidates centrally determined by MoBSE. Allocations occur annually. Although teachers are nominally required to accept their posting, in practice they may exert some control over their location by requesting a transfer

⁵ We base the information in this section on discussions and documentation from MoBSE and World Bank officials.

or refusing to accept a post. Regional administrators often honor teacher preferences by responding to such requests or refusals. Additionally, Gambia College students pursuing certification are posted to schools as teacher trainees. These trainees are in charge of their own classrooms and counted as unqualified teachers, though we explore sensitivity of results to their exclusion. Koranic teachers are posted to all government schools, but receive distinct training from other teachers, so we will also explore sensitivity of results to their exclusion.

The hardship program provides an additional allowance to teachers in lower basic grades at government schools in Regions 3-6, the regions that are farthest from the capital of Banjul and the most economically disadvantaged.⁶ Schools are classified as being in a hardship area if they are located more than 3 kilometers from a main road.⁷ The allowance is 30% of salary in Regions 3-4, 35% in Region 5, and 40% in Region 6, with the amounts intended to compensate for increasing distances from the capital Banjul. Both qualified and unqualified teachers in hardship schools receive the allowance. The allowance is funded by a World Bank-administered grant and costs an average of US\$23 per month, for an annual cost of 10.7 million GMD (US\$350,000) in 2011. The allowance is large relative to the Gambia's monthly per capita GDP of US\$43 and the average monthly teacher salary (before hardship pay) of US\$67.

By providing teachers with additional salary for working in a remote school, the hardship allowance provides monetary incentives to attract and retain teachers at these schools, which has historically been a challenge. Although teachers in hardship schools receive the allowance regardless of qualifications, the amount is larger in absolute terms for qualified teachers. Moreover, a greater willingness among qualified teachers to serve in hardship areas would give regional administrators more discretion in placing teachers. Our outcomes of interest therefore center on teacher qualifications, such as the percentage of qualified teachers in each school, and on the availability of qualified teachers to students, such as pupil-qualified teacher ratio. We are also interested in whether the hardship allowance changed the gender composition of teachers, as it was also hoped that the program would help address underrepresentation of female teachers in hardship schools. We address the impact of the hardship allowance on student enrollment and performance in a forthcoming companion paper.

⁶ Grant-aided schools are also eligible for the hardship allowance, but only two such schools are recipients among the eight in Regions 3-6. We drop these schools from the analysis in order to focus on government schools.

⁷ The definition of the main road remained constant throughout the sample period. Nonetheless, 9 schools have changed hardship status since the program's inception because their calculated distances from the main road were revised. MoBSE assigns hardship status based on distance calculations from odometer readings of vehicles driven by ministry officials. Periodic revisions to these calculations have led in some cases to hardship status changes.

3. Methodology

A. Difference-in-differences

The availability of panel data on outcomes of interest spanning time periods before and after implementation of the hardship allowance allows us to use difference-in-differences (DD) to estimate the program's impact. The difference-in-differences specification for each outcome variable, y , is:

$$y_{it} = \theta D_{it} + \delta_i + \rho_t + \gamma BCS_{it} + \varepsilon_{it} \quad (1)$$

Here, D is an indicator for treatment (hardship status) for school i in year t ; BCS is an indicator for being a Basic Cycle (combined primary and lower secondary) School; and δ and ρ are school and year fixed effects, respectively.⁸ The coefficient of interest is θ , which measures the treatment effect of hardship status on the outcome variables. We also consider a continuous treatment model, which exploits the variation in the amount of the hardship allowance, $HSPct$, across regions.

$$y_{it} = \omega HSPct_{it} + \delta_i + \rho_t + \gamma BCS_{it} + \varepsilon_{it} \quad (2)$$

The amount of the hardship allowance was 30% of teacher salary in Regions 3 and 4, 35% in Region 5, and 40% in Region 6.

In specifications (1) and (2), the impact of the hardship allowance is identified by the assumption that trends in the dependent variables over time would have been the same in hardship and non-hardship areas in the absence of the program. To investigate this assumption, we estimate the trends in the outcome variables during the years before the program was implemented. We describe the procedure in more detail and present results in Section 5.

Another potential threat to identification in the difference-in-differences analysis would be if government concurrently targeted hardship schools for policies other than the hardship allowance. In extensive discussions with MoBSE officials at both headquarters and regional offices and with school administrators, we verified that no additional policies were directed specifically to hardship schools.

B. Regression Discontinuity

The assignment of the hardship allowance based on a 3-kilometer threshold distance of the school to the main road presents an appropriate setting for a regression discontinuity (RD) design. If

⁸ In this and all subsequent equations, the constant has been omitted in favor of inclusion of a full set of relevant fixed effects, which is equivalent.

school characteristics are distributed continuously across this threshold distance and the policy was properly implemented, schools just above and just below the threshold should be nearly identical, on average, prior to receipt of the hardship allowance. Comparing these schools after receipt of the allowance would then yield an unbiased estimate of the program's impact.

The threshold distance of 3 kilometers does not perfectly predict hardship allowance receipt, however, necessitating the use of a fuzzy regression discontinuity design. We first describe the estimation method, and later address threats to its validity in the results section. The fuzzy RD design proceeds in two stages. In the first stage, we explore whether the threshold distance to the main road predicts hardship status:

$$D_i = \beta T_i + f(d_i) + \gamma BCS_i + \delta_r + \varepsilon_i \quad (3)$$

where i indexes schools; D is an indicator for treatment (hardship status); T is an indicator for being located at least 3 kilometers from the main road; $f(\cdot)$ is a flexible function of the road distance d ; BCS is an indicator for this type of school, as in (1) and (2); and δ is a region fixed effect. The coefficient of interest is β , which measures the effect of crossing the distance threshold on treatment, conditional on distance, school type, and region. We specify (3) as a linear probability model and $f(\cdot)$ as a polynomial in the running variable d .

In the second stage, we regress an outcome of interest y on the same right-hand side variables as (3) or on treatment status:

$$y_i = \rho T_i + f(d_i) + \gamma BCS_i + \delta_r + \varepsilon_i \quad (4)$$

$$y_i = \theta D_i + f(d_i) + \gamma BCS_i + \delta_r + \varepsilon_i \quad (5)$$

Equation (4) is the reduced form, measuring the effect of crossing the distance threshold on the outcome. Because no policies other than the hardship allowance are determined by crossing the distance threshold, T meets the exclusion restriction required to instrument for hardship status D . Accordingly, we estimate (5) by instrumental variables, using (3) as the first stage. The reduced-form estimate ρ measures the intent to treat (ITT). The IV estimate θ measures the local average treatment effect (LATE) for the schools induced into treatment by being just beyond the distance threshold. The LATE is equivalent to the reduced-form estimate scaled by the first stage (i.e., $\theta = \rho/\beta$).

We weight all regressions by the number of teachers in a school, because these are the recipients of the treatment. In the difference-in-differences analysis, we cluster standard errors at the school level. In the RD design, we cluster standard errors by the cluster, the sub-regional administrative units for schools, of which there are 33 in the data (we cannot cluster standard errors by school in the RD design because we work with a cross section).⁹

As the discussion of identification strategies in this section makes clear, these methods rely on different types of data (panel for DD, cross-section for RD) and different identifying assumptions (common trends between hardship and non-hardship schools for DD, balance on observable and unobservable characteristics across the distance threshold for RD). As such, they estimate different parameters. The DD estimand is the average treatment effect under an assumption of treatment effect homogeneity, while RD estimates the intent-to-treat and local average treatment effect for schools just beyond the 3km distance threshold. Relying on multiple sources of exogenous variation therefore can capture different features of the policy's effects beyond what one identification strategy would uncover alone. We return to this issue when discussing results in Section 5.C.

Another parameter that might be of interest is the average treatment effect on the treated (ATT). This parameter can be estimated using a matching-type estimator. Matching estimators require an assumption of common support, however, which clearly fails in our data when distance from a main road is included as control variable. The overlap in the support of the propensity score distribution between hardship and non-hardship schools is too sparse to generate meaningful estimates. (Note that in sharp RD designs, this overlap when matching on the running variable is zero by definition.) Thus we are limited to the assumption of a constant treatment effect when interpreting difference-in-differences estimates using the full sample.

4. Data

Data for the analysis come from several sources. Data on school personnel, student enrollment, and hardship status come from the Education Management Information System (EMIS), an annual census of schools taken by the Gambia Ministry of Basic and Secondary Education (MoBSE). We calculate the travel distance from each school to the nearest main road using school locations provided by MoBSE and a map of the road network provided by the Gambia Bureau of Statistics (GBOS).

⁹ The geographic grouping of schools into clusters and the presence of a “cluster monitor” who oversees their administration will lead to correlated errors. However, we include region rather than cluster fixed effects because the region controls teacher postings.

Additional data on pre-treatment characteristics come from the 2003 Census conducted by GBOS, for which we match schools to the nearest village. Each data source contains the universe of its respective units, and all are administrative with the exception of the Census.

In the difference-in-differences regressions below, the data include two pre-treatment years, 2001 and 2003. These years are chosen due to the availability of accurate information on teacher qualifications.¹⁰ The post-treatment years that include teacher qualifications in our data are 2010-2012.

For the regression discontinuity analysis, we use data from 2012, the most recent year available.¹¹ In all analyses, we limit the sample to government-run, Lower Basic (primary) and Basic Cycle (combined primary and lower secondary) schools in Regions 3-6. For Basic Cycle Schools, the data do not distinguish between teachers in the primary grades 1-6, who are eligible for the hardship allowance, and those in lower secondary grades 7-9, who are not. We therefore include separate intercepts for these schools in all regressions to account for any systematic differences with Lower Basic Schools.

Table 1 shows the number of schools by category over the time period of our panel. No schools were closed after the hardship program began, but several schools opened. School-level fixed effects are used to control for any systematic differences between schools in the difference-in-differences analysis. However, to guard against the possibility that entering schools have a different relationship to hardship status than schools that pre-dated the program, we repeat all DD analyses for the subsample of schools that have been open since 2001, with results nearly unchanged (results not shown but available upon request).

Table 2 presents summary statistics for the panel, separated into pre- and post-treatment periods. The dramatic changes that have recently occurred in Gambian public education are on display. During this period, school enrollment switched from being majority male to majority female; the percentage of qualified teachers nearly tripled to 64%; the pupil-qualified teacher ratio fell by 23 students; and the female qualified teacher-pupil ratio increased by fifty percent.^{12,13} To the extent that

¹⁰ Data on teacher qualifications for other pre-treatment years were estimated rather than collected directly. We do not use these data in the analysis.

¹¹ Results for 2010 and 2011 are very similar to 2012, but with a weaker first stage, because in 2012 three schools lost hardship status when they were found to be less than 3km from the main road. Results for 2012 are nearly unchanged when excluding these three schools. Results not shown but available upon request.

¹² We use the reciprocal of pupil-female qualified teacher ratio in order to include schools without female qualified teachers in the sample. The mean post-treatment qualified female teacher-pupil ratio of 0.0045 implies a pupil-female qualified teacher ratio of 222.

hardship schools were beneficiaries of these trends, a simple time-series analysis of hardship schools would risk overstating the effects of the hardship allowance, highlighting the importance of the difference-in-differences approach.

In the regression discontinuity design, the running variable is the travel distance from each school to the main road. School locations are measured with error, however, because MoBSE officials estimated school locations rather than relying on field-based collection of school coordinates. For this reason, we drop schools whose map location does not match the district (a political boundary roughly equivalent to a U.S. county) listed in the EMIS. This restriction removes 42 schools and eliminates cases in which the distance from school to road would be incorrectly measured. Remaining measurement error in school distance will attenuate our regression discontinuity estimates towards zero, giving us more confidence in findings of statistically significant effects. Finally, we drop 8 schools for which we could not find information on the nearest village in the 2003 Census. These restrictions leave a final dataset of 244 schools for our RD design. A map of the dataset appears as Figure 1 (schools in Regions 1-2 are shown on the map for illustration although they are excluded from the analysis).

The data used in the RD design contain 148 hardship and 96 non-hardship schools with 29,576 and 33,454 students, respectively. Table 3 conducts formal tests for differences between hardship and non-hardship schools across a range of characteristics. Numerous differences emerge. Hardship schools are statistically significantly less likely to be double shift (teaching two groups of students each day) and more likely to have multi-grade classrooms, reflecting their smaller size. Hardship schools are located an average of 11.8 kilometers from a main road, compared to 0.9 kilometers for non-hardship schools. They have nearly 150 fewer students and almost 5 fewer teachers, on average. They have a higher proportion of qualified teachers, 65% to 61%. However, this difference is no longer statistically significant when teacher trainees are excluded from the count of unqualified teachers. Class sizes, as measured by pupil-teacher and pupil-qualified teacher ratio, are statistically indistinguishable between hardship and non-hardship schools. Female teachers are dramatically under-represented in hardship schools, however, with a 10 percentage points lower share of the overall teaching corps and a qualified female teacher-pupil ratio that is 50% lower.

¹³ The drop in the total number of teachers between the pre- and post-treatment periods reflects a decline in the use of contract teachers, both qualified and unqualified, hired outside the civil service. Our conclusions are robust to excluding contract teachers, with results presented in Tables A3-A4 of the supplementary materials.

The bottom of Table 3 displays differences in pre-treatment characteristics, based on the 2003 Census characteristics of the school's nearest village. The relative poverty and remoteness of hardship schools emerges in these differences. Non-hardship schools are located in villages more than 3 times the size, on average, than hardship schools. They enjoy greater access to electricity, piped water and flush toilets, and have a lower illiteracy rate. The ethnic composition also differs between schools,¹⁴ and although populations surrounding both types of schools are almost entirely Muslim, even here a significant difference emerges. Taken together, the many differences between hardship and non-hardship schools and their surrounding communities presented in Table 3 highlights the inappropriateness of any empirical strategy that relies on simple comparisons between hardship and non-hardship schools in assessing the program's impact. Moreover, the likelihood that hardship and non-hardship schools also differ in unobservable characteristics make it essential to use an identification strategy that accounts for both observable and unobservable differences.

5. Results

A. Difference-in-Differences

Before presenting results of the difference-in-differences analysis, we verify that the identifying assumption of common pre-treatment trends in outcome variables between hardship and non-hardship schools holds. To investigate this assumption, we estimate trends in the outcome variables during the years before the program was implemented. Since we have two years of pre-program data on teacher characteristics, the relevant equation for those outcomes is

$$y_{it} = \tau HS_i + \delta T2_t + \beta(HS * T2)_{it} + \varphi_c + \gamma BCS_{it} + \varepsilon_{it} \quad (6)$$

Here, HS is a dummy variable indicating that a school was considered a hardship school after the program began, and $T2$ is an indicator for the year 2003. A BCS indicator and a set of cluster fixed effects, φ , are included as well. The coefficient of interest is β , which allows the impact of time to differ between the hardship and non-hardship schools.

Results are presented in Table 4. There is no evidence that hardship schools were experiencing different trends in any of the outcome variables prior to the implementation of the hardship allowance, except for the ratio of female qualified teachers to the number of students. As the identifying

¹⁴ Gambians hail from various ethnicities, with distinct traditions and languages serving as sources of identity. The national language of English is mandated in all schools, but indigenous languages are common for everyday speech.

assumption fails for this outcome, we will restrict our attention to the remaining dependent variables. We do not expect that the differential trend in the female qualified teacher-to-student ratio affected any of the other outcomes of interest. While the coefficient on the interaction, β , for this outcome is statistically significant at five percent, it is quite small in magnitude. The coefficient of 0.002 represents one additional female qualified teacher per 500 students. Given that average enrollment in hardship schools was 200 students, this difference is unlikely to be large enough to affect the decisions of other teachers or students, or to indicate any underlying differences that would invalidate identification for the other outcome equations.

The results of estimating equation (1), the main difference-in-differences specification, by OLS are presented in Table 5. The coefficient on the percentage of qualified teachers in column (1) is significant at a one-percent test size. Designation as a hardship school raises the percent of qualified teachers by 10 percentage points. This increase is substantial, given that the mean qualified teacher percentage is 45%. The effect is no longer statistically significant, however, when the denominator (total number of teachers) is restricted to exclude teacher trainees, or when both teacher trainees and Koranic teachers are excluded from the ratio, reflecting more intensive use of these types of teachers in non-hardship schools in the post-treatment period. The hardship allowance did not significantly affect the percentage of female teachers or female qualified teachers.

The hardship allowance was also not found to have a statistically significant impact on pupil to qualified teacher ratios. The coefficient on the number of qualified teachers is positive, but not significant. Enrollment, however, saw a statistically significant increase of 45 students, which is large given the average enrollment of 253 students. Thus, any possible gains in the number of qualified teachers appear to have been offset by increased enrollment, leaving the pupil to qualified teacher ratios relatively unchanged. Table 6 presents the estimates of equation (2), which specified the hardship allowance as a continuous treatment rather than binary. Once again, the only statistically significant coefficients are on the percentage of qualified teachers and enrollment. An increase of ten percentage points in the hardship allowance (the equivalent of relocating a hardship school teacher from Region 4 to 6) increases the percentage of qualified teachers by 2.7 percentage points. These results indicate that, not only has the hardship allowance succeeded in its goal of attracting qualified teachers to remote areas, but increasing the amount of the allowance has a small but measurable effect on teacher location decisions across hardship areas as well. The coefficient on enrollment suggests that the same ten

percentage point increase in the amount of the allowance would increase enrollment by 11 students, although this coefficient is only significant at 10%.

Given the large increases in the number of schools over our sample period, one might worry that if new schools placed in hardship areas disproportionately attracted teachers for reasons other than the hardship allowance (due to upgraded facilities, for instance), then our findings would be spurious. To check this, we limit the sample to schools open continuously since 2001. Doing so causes the treatment effect on qualified teacher percentage to fall from .10 to .08, but it remains significant at 5%. All other results from Tables 5-6, which we do not report but are available upon request, remain nearly unchanged.

One natural falsification exercise for our difference-in-differences specification would be to look at Upper Basic and Senior Secondary Schools before and after the program began, since teachers at these schools do not receive additional pay even if the school is located in a hardship zone. Unfortunately, the number of UBS and SSS in hardship areas was too small in the years before 2005 for this analysis to be possible.

Instead, we perform the difference-in-differences estimation on schools in Region 2. A hardship allowance of 20% was initially considered for Region 2 but not implemented, providing a useful comparison. We re-estimate equation (1), replacing hardship status with an indicator variable for being located more than three kilometers from a main road in the years after the hardship allowance. The results, which appear in Table A1 of the supplementary materials, indicate that not only is Region 2 not experiencing the same increasing percentage of qualified teachers in hardship areas as Regions 3-6, the opposite is in fact true. In Region 2, remote areas saw statistically significant decreases in the percentage (-0.21) and number (-4.3) of qualified teachers after the hardship policy was implemented in the other regions, as well as a significant increase (21.6) in the pupil-to-qualified teacher ratio. One possible explanation is that the hardship policy has made the more remote areas of Region 2 a less desirable posting for qualified teachers, who could now receive higher pay for working in the remote areas of other regions.

B. Regression Discontinuity

i. First stage

As mentioned in Section 3, the 3-kilometer distance threshold does not perfectly predict hardship status, necessitating use of a fuzzy RD design. Of the 244 schools in the analysis, 14 have hardship status that is inconsistent with their distance to the main road: 3 non-hardship schools are located more than 3 kilometers away, while 11 hardship schools are located closer than 3 kilometers. When presented with this list of apparently incorrectly assigned schools, MoBSE disputed our distance calculations and replied that their policy rule has been consistently implemented. Their own distance calculations are based on odometer readings from vehicles driven by MoBSE officials and not centrally recorded, so we cannot use their distances as an alternative measure of the running variable. Taken at face value, their response suggests that the only reason for incorrect assignment of hardship status is measurement error in our distance calculation. In extensive discussions with MoBSE officials, school administrators and teachers in all hardship regions, we did not hear any reports of successful manipulation of a school's hardship status. We examine the data for evidence of such manipulation, however, and report results below.

Table 7 presents estimates of the first stage equation (3). In column (1), we specify $f(d)$ as a quartic but do not include additional controls in the regression. The coefficient on the distance threshold T of 0.48 means that schools located just beyond the threshold are 48 percentage points more likely to be a hardship school, an effect size that is almost equivalent to the 49% of all schools that are hardship. The coefficient estimate is statistically significant at the 1% level, with an F statistic of 12.4. Columns (2)-(3) increase the polynomial order to 5 and 6, respectively, but the coefficient is almost identical. In column (4) we introduce school type and region fixed effects, and in column (5) we add a host of pre-treatment controls from the nearest village in the 2003 Census.¹⁵ Column (7) increases the polynomial order to 7. Each of these modifications leaves the coefficient on the distance threshold nearly unchanged. In column (8), we allow this 7th-order polynomial to vary on both sides of the threshold. Here we see the coefficient change by a non-negligible magnitude for the first time, but the coefficient increases to 0.81 and remains significant at 1%. The distance threshold also predicts hardship status when limiting the sample to schools in the neighborhood of the discontinuity, as in columns (8)-(9), which restrict the sample to schools within 1-5 and 2-4 kilometers from the main road, respectively.

¹⁵ The included controls are log population, employment/population ratio for ages 18+, percent with access to electricity, percent illiterate, percent Muslim, and percent of Mandinka, Fula, and Wollof ethnicities (the three largest ethnic groups).

Based on these results, our preferred specification is column (4), which offers the strongest instrument ($F=14.5$) among the specifications with a relevant first stage.¹⁶ Figure 2 presents a graphical depiction of the first stage, showing mean hardship status within bins of the running variable and the predicted hardship status from our preferred specification.

In addition to a strong first stage, valid causal inference in a regression discontinuity design requires that the running variable is not subject to manipulation. In our context, one way that such manipulation could occur is through the strategic opening or closure of schools on one side of the distance threshold. Although manipulation is plausible in either direction (those sympathetic with increasing teacher salaries would attempt to increase schools beyond the cutoff, while those interested in limiting payroll costs would attempt to increase schools inside the cutoff), the external financing of the hardship allowance makes manipulation in favor of an increase in hardship status more likely. Thus a disproportionate number of new schools just beyond the cutoff would be consistent with manipulation. In practice, no schools were closed since adoption of the hardship allowance in 2005. In that period, 16 new schools in the dataset were constructed, split equally between hardship and non-hardship. Removing these schools from the data does not change the first-stage results, as shown in column (10) of Table 7.

Manipulation of the running variable would also be likely if there is “bunching” in the density around the threshold. A histogram of distances from schools to the main road, shown in Figure 3, does not show any clear evidence of bunching. To check formally, we conduct the McCrary (2008) test for manipulation of the running variable. As mentioned above, the likelihood that any manipulation of the running variable would occur in favor of an increase in hardship schools means that the setting satisfies the monotonicity condition required of the test.¹⁷ We fail to reject the null hypothesis of no manipulation at the threshold distance of 3 kilometers (p -value=0.24), using the suggested bin width

¹⁶ Among the specifications in which the F statistic exceeds 10, column (4) is also preferred based on the Akaike and Bayesian Information Criteria (AIC and BIC). Lee and Lemieux (2010) suggest AIC as a guide to choice of polynomial order. Reduced-form second stage results allowing the polynomial to vary on both sides of the cutoff are similar to those for our preferred specification; results not shown but available upon request.

¹⁷ The required monotonicity of the McCrary (2008) test refers to the direction of manipulation of the running variable relative to its value in the absence of treatment. Note that this differs from the monotonicity required for the LATE interpretation of instrumental variables estimates obtained by (5), for which a change in the running variable must affect treatment in one direction. This latter form of monotonicity also holds in our case, because moving a school from below to above the threshold distance will increase its likelihood of receiving the hardship allowance.

and bandwidth. The result is robust to halving the reference bandwidth (p -value=0.10), as also suggested in McCrary (2008).

To verify consistent application of the hardship criterion, we also substitute straight-line distance from schools to the main road as the running variable and re-estimate (1). The coefficient on this alternative distance, shown in column (1) of Table 8, is 0.26, or nearly half the magnitude of the main first-stage estimates presented in Table 7, and significant at only the 10% level. This result helps to confirm that travel distance was the true criterion used when determining hardship status. A final check on the validity of our first stage is to look for balance of pre-treatment characteristics on either side of the distance threshold. We replace the left-hand side of (1) with a series of characteristics from the school's nearest village measured in the 2003 Census (using different variables from those included in Table 7, column [5] in order to prevent pre-test bias). The results, presented in columns (2)-(9) of Table 8, show no significant coefficients on the distance threshold, indicating that schools near the cutoff had a continuous distribution of observable characteristics prior to treatment.

ii. Second stage

The primary goal of the hardship allowance was to upgrade the quality of remotely located schools by increasing the presence of qualified teachers. Table 9 explores whether this occurred by presenting reduced-form and instrumental variables results in Panels A and B (corresponding to estimation of (4) and (5), respectively). In Panel A, column (1), the coefficient of 0.08 means that a school located just beyond the distance threshold has 8 percentage points more qualified teachers than a school located just inside the threshold. The result is not statistically significantly different from zero, however. The IV estimate in Panel B, column (1) shows that the effect of the hardship allowance is twice this magnitude (recall that IV estimates equal the reduced form divided by the first stage), but also not statistically significant. Removing teacher trainees from the count of unqualified teachers increases the magnitude of these coefficients slightly, with the IV estimate attaining significance at 10% in column (2). Removing both teacher trainees and Koranic teachers gives a similar result, with the IV estimate again marginally significant. These results show limited evidence that the hardship allowance increased teacher qualifications in hardship schools close to the distance threshold.

A secondary goal of the hardship allowance was to address other imbalances in teacher personnel in remotely located schools, including the presence of female teachers and qualified female

teachers. The hardship allowance did not significantly increase the percentage of female teachers or female qualified teachers, however, as shown in Columns (4)-(5) of Table 9.

We look for evidence that the hardship allowance increased the presence of qualified teachers on a per-student basis in column (6), using pupil-qualified teacher ratio as the outcome. The reduced-form result shows a decline of 13.2 students per qualified teacher, while the IV result is 27.4 fewer students. Both coefficients are significant at the 5% level and large relative to the mean of 45 students per qualified teacher. A graphical depiction of the reduced-form result appears in Figure 4. Column (7) shows that this result is not driven by a generalized decrease in the pupil-teacher ratio. We fail to find an increase in the female qualified teacher-pupil ratio in column (8), however. In columns (9)-(10) we unpack the pupil-qualified teacher ratio result by separating numerator and denominator (we do not weight these regressions because they are specified in levels rather than rates). Although coefficients are no longer statistically significant, the results show that the decrease in the ratio was achieved by increasing the number of qualified teachers more than proportionally to the increase in enrollment.

We have examined many outcomes, but our only statistically significant results across both reduced-form and IV estimates in the RD design is for the pupil-qualified teacher ratio. To check the robustness of this result, in Table A2 of the supplementary materials we report results of additional reduced-form regressions using pupil-qualified teacher ratio as the outcome. The decrease in pupil-qualified teacher ratio across the 3km distance threshold is robust to increasing the polynomial order to 7 and allowing this 7th-order polynomial to vary on both sides of the cutoff. The coefficient remains negative and large when using discontinuity samples of 1-5km and 2-4km from the main road, though with less statistical power. These results also survive a series of placebo tests in which we use alternate samples that should not experience a discontinuous decrease in pupil-qualified teacher ratio across the cutoff. These samples include Upper Basic (grades 7-9) and Senior Secondary (grades 10-12) schools, whose teachers must hold a different credential and do not receive the hardship allowance; Region 2, which was considered for a 20% hardship allowance that was ultimately not implemented; and data from 2001 and 2003, before the program was in place. The coefficient on the distance threshold is not statistically significant in any of these samples.

As a final check on the plausibility of our pupil-qualified teacher results, we follow Card, Mas, and Rothstein (2008) and Kazianga et al. (2013) and check for alternative distance thresholds by modifying (4) as:

$$PQTR_i = \rho \mathbf{1}(d_i \geq a) + \gamma BCS_i + \delta_r + \varepsilon_i \quad (7)$$

where $PQTR$ is pupil-qualified teacher ratio and $\mathbf{1}(\cdot)$ is the indicator function. We vary the choice of a along a grid from [0.1,10] in increments of 0.1 kilometers. Hansen (2000) shows that in this procedure, the best fitting equation among alternative thresholds consistently estimates the threshold a . Figure 5 presents results, plotting the R^2 from each regression as a function of the threshold a . The threshold of 3 kilometers gives the best fit to the data, consistent with the hardship allowance causing a decrease in the pupil-qualified teacher ratio.

C. Discussion

We have evaluated the impact of the hardship allowance on the distribution of teachers across schools using two different methodologies. Using a difference-in-differences approach, the only outcome with a robustly significant impact is the percentage of qualified teachers. In the regression discontinuity design, the program appears to have affected only the pupil-qualified teacher ratio. Are these results contradictory, or do they paint a consistent picture of the program impact?

In comparing results between the two approaches, it is useful to recall that the estimators apply to different subpopulations. While the difference-in-differences estimates compare hardship and non-hardship schools in the full sample, the regression discontinuity results apply only to schools near the distance threshold. If the hardship allowance affected schools near the cutoff differently than the rest of the sample, estimates obtained from the two approaches should reflect these differences.

Consider the results for the percentage of qualified teachers. Our difference-in-difference estimate is an increase in qualified teachers of 10 percentage points due to the hardship allowance (Table 5, column [1]). Our IV estimate in the regression discontinuity design is 0.16 (Table 9, column [1]), which is larger in magnitude but not significant at conventional levels. Restricting the sample to Lower Basic schools only might provide a sharper estimate, because in Basic Cycle schools we cannot distinguish teachers in primary grades (who are eligible for hardship pay) from lower secondary grades (who are not). These results, shown in Table 10, Panels A-B reveal that the hardship impact on qualified teacher percentage is more than twice the magnitude in the RD design than difference-in-differences, with both estimates statistically significant.

This difference is consistent with a scenario in which teachers are more willing to serve in hardship schools that are relatively close to main roads, but more reluctant to teach as this distance

increases. Because the hardship allowance does not vary with distance once a school crosses the 3 kilometer threshold, this interpretation seems plausible. It is also consistent with anecdotal reports from regional education officers about difficulties in placing qualified teachers in the most remotely located schools. Re-estimating the difference-in-differences regressions on samples close to the distance threshold, as in Table 10, Panels C-D provides further confirmation of this explanation. The coefficients on qualified teacher percentage generally increase in magnitude relative to the full sample and approach those of the RD design.

In the case of the pupil-qualified teacher ratio, a similar phenomenon likely applies. In that case, our point estimates indicate increases in the number of qualified teachers and student enrollment in response to the hardship allowance in both the difference-in-differences and regression discontinuity approaches (Tables 5 and 9, columns [9]-[10]). It is therefore the relative magnitude of these increases, not their direction, which accounts for differences in impact estimates for pupil-qualified teacher ratio in the two approaches. If teachers respond strongly to the hardship allowance close to the 3 kilometer cutoff but students do not, we would obtain this result. This interpretation seems plausible, because teachers operate in a national labor market, whereas students are typically constrained to attend their village school. As we would expect, both our RD and DD estimates increase in magnitude when restricted to discontinuity samples close to the threshold (Table A4, columns [3]-[4]; and Table 10, Panels C-D, column [4]).

D. Substitution or Scale Effect?

Given the evidence we have found in favor of the hardship allowance successfully increasing the presence of qualified teachers in hardship schools, it is natural to wonder whether this was the result of qualified teachers switching from non-hardship to hardship schools, or whether the allowance led to more teachers becoming qualified. In other words, did the program work through a substitution or scale effect?

A full decomposition of substitution and scale effects in response to the allowance would require detailed data on teacher certification dates and movements between schools. Such data are unavailable. Instead, we conduct a back-of-the-envelope calculation to gauge whether the likely increase in qualified teacher labor supply due to the hardship allowance was sufficient to meet the increased demand in hardship schools. The exercise proceeds in two steps. First, we calculate the magnitude of the increase in qualified teacher demand implied by our regression estimates. Next, we

use data on teacher certifications to estimate the increase in qualified teacher supply due to the hardship allowance. If the estimated increase in demand exceeds supply, then the difference would represent a lower bound on the substitution of qualified teachers between non-hardship and hardship schools. If the increased supply exceeds demand, then substitution would not have been necessary.

First, we find the increase in demand for qualified teachers in hardship schools by multiplying the estimated treatment effect for qualified teacher percentage (Table 5, column [1]) by the average number of teachers in hardship schools in the post-treatment period: $0.10 \times 1,401 = 140$.

We next estimate the increase in supply of qualified teachers due to the hardship allowance. In the Gambia, teachers become certified by completing a three-year training course at the Gambia College School of Education. Those who complete the Primary Teaching Certificate (PTC) are qualified to teach primary grades, while those who complete the Higher Teaching Certificate (HTC) are qualified to teach lower secondary grades. Importantly, HTC teachers are not posted to primary grades and therefore are ineligible for the hardship allowance.

Figure 6 shows trends in teacher certifications using data from the Gambia College;²⁰ the year 2007 divides the pre- and post-treatment period because those graduating in 2007 would have entered training in 2005, when the hardship allowance began. The graph suggests common pre-treatment trends between PTC (primary) and HTC (lower secondary) certifications; if anything, primary certifications appear to be trending downward in the years immediately before treatment. In the post-treatment period, primary certifications show a clear upward trend, while lower secondary certifications grow only modestly. The simple difference-in-differences estimate from this data yields 310 additional certified primary school teachers, on average, in the post-treatment period. This amount represents an annual *flow* of new qualified teachers available for primary schools, greatly exceeding the *stock* of 140 added to hardship schools.²¹

In practice, the increased demand for qualified teachers in hardship schools was likely met by a combination of new entrants, previously unqualified teachers who became certified, and teachers who

²⁰ Data provided directly by Gambia College for 2005-2012 are augmented with Gambia College data for 2002-2004 reported in Mulkeen (2007). Removing the 2002-2004 data and repeating the exercise in this section leads to nearly identical results.

²¹ Several elements of the certification environment changed differentially for prospective primary and lower secondary trainees in the post-treatment period, other than the hardship allowance. Using a conservative set of assumptions to adjust our estimates for these changes still yields a sufficient flow of new qualified teachers into hardship schools to meet demand. Details are available in the working paper version (Pugatch and Schroeder 2013).

switched from non-hardship schools. Unfortunately, available data do not allow us to distinguish these channels directly. However, our estimates of the relative increases in supply and demand suggest that substitution of qualified teachers from non-hardship schools was not necessary to achieve the increases observed in the data.

6. Conclusion

We have evaluated the impact of the Gambian hardship allowance on the distribution of qualified teachers across schools, using the exogenous variation generated by common pre-treatment trends between hardship and non-hardship schools and the geographic discontinuity in treatment. We find that the hardship allowance increased the presence of qualified teachers by 10 percentage points in hardship schools. A comparison of this difference-in-differences estimate with regression discontinuity results suggests that the gains were larger in hardship schools nearer to the 3-kilometer distance threshold. We also find a decrease in the pupil-qualified teacher ratio of 27, although these gains are localized to those hardship schools closest to the distance threshold. These findings suggest that the hardship allowance succeeded in its primary goal of recruiting qualified teachers to rural areas, but that it was less successful in the most remotely located schools. We also do not find evidence that the allowance succeeded in addressing the low share of female teachers in remote schools.

From a policy perspective, the hardship allowance would be less effective if it merely reallocated qualified teachers from less remote to more remote schools. The country's schools will be better served if the policy succeeds in attracting more qualified teachers to the system rather than spreading existing qualified teachers more evenly. Although we cannot formally decompose these channels using available data, a back-of-the-envelope calculation suggests that the increased flow of qualified teachers after the adoption of the policy was sufficient to meet the increased demand.

Our back-of-the-envelope calculation also revealed that the hardship allowance added 140 more qualified teachers to hardship schools beyond what they would have experienced in the absence of the policy. At a cost of approximately US\$350,000 annually, these additional teachers cost US\$2,500 each to recruit, a considerable expense in a country with a per capita GDP of US\$512. Of course, the hardship allowance may also bring additional benefits, such as increased morale, to all recipients, qualified or unqualified, or stimulus to village economies through teachers' increased purchasing power. Ultimately its success hinges on whether it improves outcomes for students, a topic we will address in forthcoming work. However, the narrower goal of attracting qualified teachers to remote schools may be achieved

more cost effectively by an amended policy. Increasing the gap in the allowance between qualified and unqualified teachers, for instance, or increasing the allowance for the most remotely located schools might better orient program resources towards recruitment goals.

Acknowledgements

We thank officials at the Gambia Ministry of Basic and Secondary Education, Gambia Bureau of Statistics, Gambia College School of Education, Gambian office of the West African Examinations Council, and the World Bank, particularly Alpha Bah, Momodou Cham, Jenny Hsieh, Sherif Yunus Hydara, Nathalie Lahire, Palamin Mbowe, Aidan Mulkeen, Isatou Ndow, Momodou Sanneh, Ryoko Tomita, and Bassirou Touré. We thank Marc Bellemare, Alfredo Burlando, Patrick Emerson, Osborne Jackson, Laura Kawano, Carol Tremblay, and Vic Tremblay for helpful comments. We are especially grateful to Candice-Michelle Weems for GIS assistance. Emily Edwards and Andrew Spaeth provided excellent research assistance.

References

- Antos, Joseph R., and Sherwin Rosen. 1975. "Discrimination in the Market for Public School Teachers." *Journal of Econometrics* 3 (2): 123–50.
- Banerjee, Abhijit V., Shawn Cole, Esther Duflo, and Leigh Linden. 2007. "Remedying Education: Evidence from Two Randomized Experiments in India." *Quarterly Journal of Economics* 122 (3): 1235–64.
- Barbieri, Gianna, Claudio Rossetti, and Paolo Sestito. 2011. "The Determinants of Teacher Mobility: Evidence Using Italian Teachers' Transfer Applications." *Economics of Education Review* 30 (6): 1430–44.
- Benabou, Roland, Francis Kramarz, and Corinne Prost. 2009. "The French Zones d'Education Prioritaire: Much Ado about Nothing?" *Economics of Education Review* 28 (3): 345–56.
- Boyd, Donald, Hamilton Lankford, Susanna Loeb, and James Wyckoff. 2005. "The Draw of Home: How Teachers' Preferences for Proximity Disadvantage Urban Schools." *Journal of Policy Analysis and Management* 24 (1): 113–32.
- . 2013. "Analyzing the Determinants of the Matching of Public School Teachers to Jobs: Disentangling the Preferences of Teachers and Employers." *Journal of Labor Economics* 31 (1): 83–117.
- Card, David, Alexandre Mas, and Jesse Rothstein. 2008. "Tipping and the Dynamics of Segregation." *Quarterly Journal of Economics* 123 (1): 177–218.
- Clotfelter, Charles, Elizabeth Glennie, Helen Ladd, and Jacob Vigdor. 2008. "Would Higher Salaries Keep Teachers in High-Poverty Schools? Evidence from a Policy Intervention in North Carolina." *Journal of Public Economics* 92 (5-6): 1352–70.
- Darling-Hammond, Linda, Barnett Berry, and Amy Thoreson. 2001. "Does Teacher Certification Matter? Evaluating the Evidence." *Educational Evaluation and Policy Analysis* 23 (1): 57–77.
- Duflo, Esther, Pascaline Dupas, and Michael Kremer. 2011. "Peer Effects, Teacher Incentives, and the Impact of Tracking: Evidence from a Randomized Evaluation in Kenya." *American Economic Review* 101 (5): 1739–74.
- Galiani, Sebastian, and Ricardo Perez-Truglia. 2011. "School Management in Developing Countries". SSRN Scholarly Paper ID 1972459. Rochester, NY: Social Science Research Network.
- 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". Working Paper 17554. National Bureau of Economic Research.
- Hansen, Bruce E. 2000. "Sample Splitting and Threshold Estimation." *Econometrica* 68 (3): 575–603.
- Hanushek, Eric A. 2003. "The Failure of Input-Based Schooling Policies." *Economic Journal* 113 (485): F64–98.
- Hanushek, Eric A., John F. Kain, and Steven G. Rivkin. 2004. "Why Public Schools Lose Teachers." *The Journal of Human Resources* 39 (2): 326–54.
- Hough, Heather J. 2012. "Salary Incentives and Teacher Quality: The Effect of a District-Level Salary Increase on Teacher Recruitment."
- Jaramillo, Miguel. 2012. "The Spatial Geography of Teacher Labor Markets: Evidence from a Developing Country." *Economics of Education Review* 31 (6): 984–95.
- Kazianga, Harounan, Dan Levy, Leigh L. Linden, and Matt Sloan. 2013. "The Effects of 'Girl-Friendly' Schools: Evidence from the BRIGHT School Construction Program in Burkina Faso." *American Economic Journal: Applied Economics* 5 (3): 41–62.
- Lee, David S., and Thomas Lemieux. 2010. "Regression Discontinuity Designs in Economics." *Journal of Economic Literature* 48 (2): 281–355.
- McCrary, Justin. 2008. "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test." *Journal of Econometrics* 142 (2): 698–714.

- McEwan, Patrick J. 1999. "Recruitment of Rural Teachers in Developing Countries: An Economic Analysis." *Teaching and Teacher Education* 15 (8): 849–59.
- . 2013. "Improving Learning in Primary Schools of Developing Countries: A Meta-Analysis of Randomized Experiments." *Unpublished Manuscript, Wellesley College*.
- Mulkeen, Aidan. 2007. "Teacher Issues in The Gambia." Workshop Background Paper.
- Mulkeen, Aidan, and Dandan Chen. 2008. *Teachers for Rural Schools: Experiences in Lesotho, Malawi, Mozambique, Tanzania, and Uganda*. Africa Human Development Series.
- Muralidharan, Karthik, and Venkatesh Sundararaman. 2013. "Contract Teachers: Experimental Evidence from India". Working Paper 19440. National Bureau of Economic Research.
<http://www.nber.org/papers/w19440>.
- Pugatch, Todd, and Elizabeth Schroeder. 2013. "Incentives for Teacher Relocation: Evidence from the Gambian Hardship Allowance". IZA Discussion Paper 7723. Institute for the Study of Labor (IZA).
- Rivkin, Steven G., Eric A. Hanushek, and John F. Kain. 2005. "Teachers, Schools, and Academic Achievement." *Econometrica* 73 (2): 417–58.
- Urquiola, Miguel, and Emiliana Vegas. 2005. "Arbitrary Variation in Teacher Salaries: An Analysis of Teacher Pay in Bolivia." In *Incentives to Improve Teaching: Lessons from Latin America*, edited by Emiliana Vegas, 187–211. Washington, D.C.:

Figure 1: Gambia Hardship Allowance, 2012




Figure 2: Regression Discontinuity, first stage




Figure shows mean of hardship status within bins defined by distance from main road (bandwidth=0.2km). Line is predicted hardship status from first stage regression, as in column (4) of Table 3.

Figure 3: Density of distance from schools to main road




Figure 4: Pupil-qualified teacher ratio




Figure 5: Pupil-qualified teacher ratio, alternative thresholds




Figure 6: Teacher Certifications




Table 1: Number of schools by type and status over time

	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012
LBS	175	211	211	207	209	223	230	228	227	231	236	236	247	250	254
BCS	2	2	4	14	14	22	30	37	38	38	38	38	40	38	40
Non-hardship								98	98	98	100	100	106	107	112
Hardship								167	167	171	174	174	181	181	182
Total	177	213	215	221	223	245	260	265	265	269	274	274	287	288	294

Table reports count of schools in each cell. Dataset is government-run Lower Basic and Basic Cycle schools in Regions 3-6. Source: EMIS.

Table 2: Summary statistics, school panel

	Pre-treatment period (2001, 2003)		Post-treatment period (2010, 2011, 2012)	
	n=322		n=853	
	mean	std. dev.	mean	std. dev.
enrollment	250.1	199.3	254.4	225.0
male enrollment	128.0	107.2	121.3	109.6
female enrollment	121.2	96.9	133.0	118.1
number of teachers	18.5	14.7	8.3	6.2
qualified teachers	4.0	4.6	5.3	4.6
teacher trainees	1.1	1.9	1.7	2.1
% qualified teachers*	0.22	0.09	0.64	0.18
% qualified teachers (exclude TT)*	0.23	0.10	0.80	0.19
% female teachers*	0.16	0.14	0.22	0.15
pupil-teacher ratio*	13.5	7.15	30.6	10.0
pupil-qualified teacher ratio*	75.7	53.3	52.4	24.3
female qualified teacher-pupil ratio*	0.0031	0.0047	0.0045	0.0045

Table reports statistics by school-year (n). Dataset is government-run Lower Basic and Basic Cycle schools in Regions 3-6. Source: EMIS.

* weighted by number of teachers.

Table 3: Balance tests, 2012

Variable	Non-hardship	Hardship	Difference
No. of schools	96	148	52
No. of students	33,454	29,576	-3,878
<u>school characteristics</u>			
double shift	0.58	0.47	-0.11*
multi-grade	0.18	0.45	0.27***
distance from main road (km)	0.9	11.8	10.9***
<u>enrollment</u>			
total	348.5	199.8	-148.6***
male	166.1	95.9	-70.2***
female	182.4	104.0	-78.5***
<u>teachers</u>			
total	12.6	7.8	-4.8***
qualified	7.6	5.1	-2.5***
trainees	2.9	1.5	-1.4***
unqualified	2.1	1.2	-0.9***
% qualified	0.61	0.65	0.05**
% qualified (excl. trainees)	0.80	0.82	0.02
% female	0.26	0.16	-0.1***
pupil-teacher ratio	27.4	25.7	-1.7
pupil-qualified teacher ratio	46.9	42.9	-4.1
qualified female teacher-pupil ratio	0.006	0.004	-0.002***
<u>village characteristics (2003)</u>			
population	7,311	2,338	-4,973***
e/pop ratio, age 18+	0.66	0.76	0.1***
access to electricity (%)	0.16	0.02	-0.14***
access to piped water (%)	0.11	0.02	-0.09***
access to flush toilet (%)	0.03	0.01	-0.02***
illiteracy	0.28	0.36	0.08***
Muslim	0.99	1.00	0.01***
Mandinka	0.44	0.31	-0.13***
Fula	0.20	0.19	-0.01
Wollof	0.20	0.13	-0.08**
married	0.52	0.56	0.04***
polygamous (conditional on married)	0.47	0.53	0.06***

Table shows mean school-level characteristics by hardship status, with village characteristics found by matching school to nearest village. Final column reports difference in means (hardship minus non-hardship), with significance levels 10%, 5% and 1% indicated by 1, 2, and 3 stars, respectively (not applicable for number of schools or students). Teacher characteristics weighed by number of teachers in school, while village characteristics weighted by population. Sample limited to government-run Lower Basic and Basic Cycle schools in Regions 3-6 only. Sample excludes schools whose map location does not match district reported in EMIS. Sources: 2003 Census, 2012 EMIS.

Table 4: Pre-treatment trends

Dependent variable	QT%	QT%	QT%	female	female	pupil-	pupil-teacher	female QT -	qualified	enrollment
		(excl. TTs)	(excl. TTs & KTs)	teacher %	QT%	QT	ratio	pupil ratio	teachers	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
hardship*time	0.02	0.02	0.02	0.03	0.02	-5.0	1.2	0.002**	0.8	41.9
	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)	(5.7)	(0.9)	(0.001)	(0.8)	(35.1)
time	-0.01	-0.02	-0.01	-0.02	-0.02	3.7	-0.2	-0.0013**	-1.73***	-54.9*
	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)	(4.2)	(0.6)	(0.0005)	(0.66)	(29.0)
hardship	-0.04**	-0.06***	-0.06***	-0.09***	-0.07*	0.2	-2.0*	-0.003**	-2.54***	-160.2***
	(0.02)	(0.02)	(0.02)	(0.02)	(0.04)	(7.1)	(1.1)	(0.001)	(0.81)	(38.8)
Mean of dependent variable	0.22	0.23	0.23	0.16	0.14	75.7	13.5	0.0031	4.0	250.1
Observations	322	322	322	322	322	322	322	322	322	322
R-squared	0.39	0.44	0.44	0.46	0.30	0.49	0.31	0.40	0.48	0.4

Robust standard errors in parentheses, clustered by school. *** p<0.01, ** p<0.05, * p<0.1. Sample is government-run, Lower Basic and Basic Cycle Schools in Regions 3-6. Data are from pre-treatment period only: 2001 and 2003 if outcome relies on teacher qualifications, and 1998-2004 otherwise. Regressions include a school type dummy and school-level fixed effects. All regressions weighted by number of teachers employed at school, except columns (9)-(10), which are unweighted. Source: EMIS.

Table 5: Difference-in-differences results: hardship status

Dependent variable	QT%	QT%	QT%	female	female	pupil-QT	pupil-teacher	female QT -	qualified	enrollment
		(excl. TTs)	(excl. TTs & KTs)	teacher %	QT%	ratio	ratio	pupil ratio	teachers	
hardship status	0.10***	0.05	0.05	-0.02	-0.04	-0.2	0.9	0.0001	0.9	45.2**
	(0.03)	(0.04)	(0.04)	(0.03)	(0.04)	(9.3)	(1.9)	(0.001)	(0.7)	(20.3)
Mean of dependent variable	0.45	0.54	0.54	0.19	0.17	63.0	22.8	0.0039	4.9	253.0
Observations	1,174	1,174	1,174	1,174	1,174	1,174	1,174	1,174	1,174	1,174
R-squared	0.84	0.89	0.88	0.67	0.63	0.59	0.71	0.57	0.84	0.94

Robust standard errors in parentheses, clustered by school. *** p<0.01, ** p<0.05, * p<0.1. Sample is government-run, Lower Basic and Basic Cycle Schools in Regions 3-6. Data are from 2001, 2003, and 2010-2012 if outcome relies on teacher qualifications, and 1998-2012 otherwise. Regressions include time dummies, a school type dummy, and school-level fixed effects. All regressions weighted by number of teachers employed at school, except columns (9)-(10), which are unweighted. Source: EMIS.

Table 6: Difference-in-differences results: amount of hardship allowance

Dependent variable	QT%	QT%	QT%	female	female	pupil-QT	pupil-teacher	female QT -	qualified	enrollment
		(excl. TTs)	(excl. TTs & KTs)	teacher %	QT%	ratio	ratio	pupil ratio	teachers	
% hardship allowance	0.27*** (0.10)	0.15 (0.10)	0.13 (0.11)	-0.07 (0.09)	-0.12 (0.12)	-13.7 (27.7)	1.7 (5.3)	0.0003 (0.004)	2.6 (1.8)	112.2* (57.2)
Mean of dependent variable	0.45	0.54	0.54	0.19	0.17	63.0	22.8	0.0039	4.9	253.0
Observations	1,174	1,174	1,174	1,174	1,174	1,174	1,174	1,174	1,174	1,174
R-squared	0.84	0.89	0.88	0.67	0.63	0.59	0.71	0.57	0.84	0.94

Robust standard errors in parentheses, clustered by school. *** p<0.01, ** p<0.05, * p<0.1. Sample is government-run, Lower Basic and Basic Cycle Schools in Regions 3-6. Data are from 2001, 2003, and 2010-2012 if outcome relies on teacher qualifications, and 1998-2012 otherwise. Regressions include time dummies, a school type dummy, and school-level fixed effects. All regressions weighted by number of teachers employed at school, except columns (9)-(10), which are unweighted. Source: EMIS.

Table 7: Regression Discontinuity, Stage 1 results

	<u>full sample</u>							<u>discontinuity samples</u>		<u>excludes schools</u>
								<u>1-5km</u>	<u>2-4km</u>	<u>opened since 2005</u>
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Distance ≥ 3km	0.48	0.47	0.48	0.48	0.47	0.48	0.81	0.69	1.16	0.48
	(0.14)***	(0.14)***	(0.13)***	(0.13)***	(0.13)***	(0.13)***	(0.27)***	(0.25)***	(0.46)**	(0.13)***
Observations	244	244	244	244	244	244	244	69	33	228
R-squared	0.87	0.88	0.88	0.88	0.89	0.88	0.9	0.65	0.53	0.89
F-stat on distance cutoff	12.4	11.4	13.4	14.5	13.9	14.4	8.8	7.8	6.3	12.6
Mean of dependent variable	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.58	0.69	0.49
Polynomial order	4	5	6	6	6	7	7	6	6	6
Region and school type fixed effects				x	x	x	x	x	x	x
Controls included					x					
Polynomial varies on either side of cutoff							x			

Regressions are linear probability models of school hardship status on travel distance from school to main road. I.e., regressions are Stage 1 of fuzzy RD design for treatment of hardship allowance. Robust standard errors in parentheses, clustered by cluster (sub-regional school administrative unit, of which there are 33 in sample). * significant at 10%; ** significant at 5%; *** significant at 1%. Sample is government-run Lower Basic and Basic Cycle Schools, Regions 3-6, in 2012. Sample excludes schools whose map location does not match district reported in EMIS. All regressions weighted by number of teachers employed at school. All regressions include polynomial in distance of indicated order. Regression controls included where indicated, from 2003 Census data on nearest settlement to school: log population, employment/population ages 18+, percent with access to electricity, percent illiterate, percent Muslim, percent of Mandinka, Fula, and Wolof ethnicities.

Table 8: Regression Discontinuity, Stage 1 robustness checks

	<u>straight-line distance</u>			<u>pre-treatment village characteristics</u>					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	hardship allowance	schooling above primary	access to piped water	access to flush toilet	marriage rate	polygamous marriage	born outside village	access to television	children born
Distance \geq 3km	0.26 (0.14)*	0.04 (0.06)	0.03 (0.02)	-0.01 (0.01)	-0.01 (0.03)	0.04 (0.05)	0.02 (0.03)	-0.04 (0.09)	-0.14 (0.24)
Observations	244	244	244	244	244	244	244	244	244
R-squared	0.86	0.22	0.17	0.06	0.12	0.15	0.25	0.19	0.16
F-stat on distance cutoff	3.5	0.5	2.8	0.7	0.2	0.6	0.5	0.2	0.3
Mean of dependent variable	0.49	0.33	0.03	0.01	0.55	0.46	0.14	0.55	3.06

Table shows results of regressions of school hardship status (column 1) or village characteristics from 2003 Census on indicator for school distance more than 3km from main road. All regressions follow preferred Stage 1 specification and include 6th-order polynomial in distance, and region and school type fixed effects. Sample is government-run Lower Basic and Basic Cycle Schools, Regions 3-6, in 2012. Sample excludes schools whose map location does not match district reported in EMIS. Schools matched to nearest village to assign village characteristics. Column (1) replaces distance along road with straight-line distance. Column (9) dependent variable is average number of children born to women ages 12 and older. All regressions weighted by number of teachers at school. Robust standard errors in parentheses, clustered by cluster (sub-regional school administrative unit, of which there are 33 in sample). * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 9: Regression Discontinuity, Stage 2 results

Dependent variable	<u>QT%</u>	<u>QT%</u>	<u>QT%</u>	<u>female</u>	<u>female QT%</u>	<u>pupil-QT</u>	<u>pupil-teacher</u>	<u>female QT-</u>	<u>qualified</u>	<u>enrollment</u>
		(excl. TTs)	(excl. TTs & KTs)	teacher %		ratio	ratio	pupil ratio	teachers	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
<u>Panel A: reduced form</u>										
Distance ≥ 3km	0.08	0.09	0.10	0.04	0.02	-13.2	-3.9	0.001	1.9	37.5
	(0.05)	(0.06)	(0.06)	(0.04)	(0.06)	(5.5)**	(4.6)	(0.002)	(1.0)*	(32.9)
R-squared	0.05	0.1	0.1	0.25	0.17	0.17	0.19	0.2	0.27	0.3
<u>Panel B: Instrumental variables</u>										
hardship allowance	0.16	0.19	0.20	0.08	0.04	-27.4	5.2	0.002	5.2	100.8
	(0.11)	(0.10)*	(0.10)*	(0.09)	(0.12)	(12.7)**	(3.4)	(0.003)	(3.4)	(104.1)
Observations	244	244	244	244	244	244	244	244	244	244
1st stage F-stat	14.5	14.5	14.5	14.5	14.5	14.5	14.5	14.5	7.6	7.6
Mean of dependent variable	0.63	0.81	0.8	0.21	0.2	45.0	26.6	0.005	6.1	258.3

Table shows results of regressions of school outcomes on distance threshold or hardship allowance receipt, as indicated. Panel B uses distance threshold to instrument for hardship allowance. Dependent variables are qualified teacher percentage (QT%), excluding teacher trainees (TTs) and Koranic teachers (KTs) as indicated; female teacher percentage and female QT% (i.e., female QTs/total QTs); pupil-qualified teacher, pupil-teacher and female qualified teacher-pupil ratios; and counts of qualified teachers and enrolled students. All regressions include a 6th-order polynomial in distance from school to main road, and region and school type fixed effects. All regressions weighted by number of teachers in school, except columns (9)-(10), which are unweighted. Robust standard errors in parentheses, clustered by cluster (sub-regional school administrative unit, of which there are 33 in sample). * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 10: Regressions using Lower Basic Schools and discontinuity samples only

Dependent variable	QT% (1)	QT% (excl. TTs) (2)	QT% (excl. TTs & KTs) (3)	pupil-QT ratio (4)
<u>Panel A: difference-in-differences, LBS only</u>				
hardship status	0.08** (0.03)	0.03 (0.04)	0.03 (0.04)	0.72 (9.34)
Observations	1,038	1,038	1,038	1,038
R-squared	0.85	0.89	0.87	0.66
Mean of dependent variable	0.50	0.61	0.60	66.8
<u>Panel B: regression discontinuity (IV), LBS only</u>				
hardship status	0.22** (0.10)	0.21** (0.11)	0.22** (0.11)	-22.1** (10.8)
Observations	209	209	209	209
First stage F-stat	16.8	16.8	16.8	16.8
Mean of dependent variable	0.63	0.81	0.80	44.9
<u>Panel C: difference-in-differences, discontinuity sample (1-5km from main road)</u>				
hardship status	0.09* (0.06)	0.18* (0.10)	0.10* (0.06)	-2.8 (15.0)
Observations	296	296	296	296
R-squared	0.82	0.80	0.88	0.61
Mean of dependent variable	0.52	0.64	0.64	59.1
<u>Panel D: difference-in-differences, discontinuity sample (2-4km from main road)</u>				
hardship status	0.14* (0.08)	0.21* (0.12)	0.07 (0.13)	-20.3 (19.5)
Observations	141	141	141	141
R-squared	0.80	0.88	0.89	0.59
Mean of dependent variable	0.57	0.69	0.68	55.1

Table shows results of regressions of school outcomes on hardship allowance receipt, as indicated. Sample is government-run Lower Basic and Basic Cycle Schools in Regions 3-6. Panels A-B drop Basic Cycle Schools. Panels C-D restrict sample to schools within indicated distance of main road. Sample in Panel B excludes schools whose map location does not match district reported in EMIS. Panels A, C and D are difference-in-differences specifications, using data from 2001, 2003, and 2010-2012. Panel B is instrumental variables estimates from regression discontinuity design, where distance threshold instruments for hardship allowance. All regressions include school type fixed effects. Panels A, C and D regressions include school fixed effects and time dummies. Panel B regressions include region fixed effects and a 6th-order polynomial in distance from school to main road. All regressions weighted by number of teachers in school. Robust standard errors in parentheses, clustered by school in Panels A, C and D and cluster (sub-regional school administrative unit, of which there are 33 in sample) in Panel B. * significant at 10%; ** significant at 5%; *** significant at 1%. Source: EMIS.