Teacher Compensation and Structural Inequality:

Evidence from Centralized Teacher School Choice in Peru[†]

Matteo Bobba Tim Ederer Gianmarco León-Ciliotta
Christopher A. Neilson Marco Nieddu

March 29, 2021

Abstract

This paper studies how increasing teacher compensation at hard-to-staff schools can reduce structural inequality in the access to high-quality teachers. We first document dramatic inequities in schooling inputs and teacher quality to which students have access in the context of a large and diverse developing country: Perú. We then leverage a change in teacher compensation to show causal evidence that increasing salaries at less desirable public schools attracts better quality applicants and improves subsequent student test scores. We finally estimate a model of teacher preferences over local community amenities, school characteristics and wages using detailed job posting and application data from the country-wide centralized teacher assignment system. The fitted model is able to replicate the main features in the data, including the sorting patterns of teachers around the policy change in teacher wages. Model estimates indicate that while current pay bonuses in less desirable regions are helpful, the current policy is woefully insufficient to compensate teachers for the lack of school and community amenities, especially in school vacancies that are distant to the teachers' home town or the location of their current job. Counterfactual experiments taking into account equilibrium sorting show that budget-neutral changes in the current wage schedule can achieve a remarkably more equitable distribution of teacher quality across regions.

Keywords: Inequality, Teacher Wages, Teacher Quality, Student Achievement.
JEL Codes: J31, J45, I21, C93, O15.

[†]We thank the Ministry of Education of Perú and in particular Jose Maria Rentería and Mariella Pérez Costa for their interest in supporting rigorous policy evaluation efforts and continued collaboration. Bobba thanks financial support from the AFD, the H2020-MSCA-RISE project GEMCLIME-2020 GA No 681228, and the ANR under grant ANR-17-EURE-0010 (Investissements d'Avenir program). León-Ciliotta recognizes financial support from the Spanish Ministry of Economy and Competitiveness (through the Severo Ochoa Programme for Centres of Excellence in R&D (SEV-2011-0075), grant ECO2014-55555-P and the Ramon y Cajal Fellowship RyC-2017-23172) and the Barcelona GSE Seed Grant. Matteo Bobba and Tim Ederer are affiliated with Toulouse School of Economics, University of Toulouse Capitole. Gianmarco León-Ciliotta is affiliated with Universitat Pompeu Fabra & Barcelona GSE & IPEG & CEPR. Christopher A. Neilson is affiliated with Princeton University. Marco Nieddu is affiliated with the University of Cagliari & CRENoS. The first version of this paper was titled *Teacher Wages, Student Achievement and the Recruitment of Talent in Rural Peru* and was presented in 2017.

1 Introduction

Children born in poorer and more rural communities face significant disadvantages in their ability to invest in their human capital. Some of these disadvantages are the product of the inequities from the past. Current policies can also reinforce inequality that is structural by providing unequal access to public school funding, investment in infrastructure and quality of instruction. This can create a feedback loop where poor education hampers local economic growth which makes it more difficult to adequately fund the local public-school system. This paper studies how teacher compensation for public school teaching jobs can contribute to reducing the structural inequality in education in Peru.

The level and structure of public sector compensation play a key role in the ability of governments to attract, retain and motivate high-quality employees. However, contracts in the public sector typically feature quite rigid wage profiles, often exclusively based on seniority, which lead to workers sorting on non-pecuniary aspects of employment (Rosen 1986). This issue is particularly important for the provision of services in jobs or locations where working conditions are less appealing, which therefore attract low-quality applicants or none at all. In the education sector, this can translate into large and persistent differences in teacher quality across communities in locations with varying levels of amenities valued by teachers.[1] In this paper we argue that given teacher preferences for community and school amenities, rigid teacher compensation profiles contribute to the structural inequities faced by children born in poorer rural regions of Peru. We also show evidence that reforming teacher pay to compensate for a lack of local amenities can help to reduce structural inequality.

We begin by documenting that a child will receive strikingly different school inputs depending upon where she is born. As in many countries, an important challenge in Peru is to provide access to quality education in rural areas. When compared to urban schools, students in rural schools are taught by teachers with lower competency scores, and are less likely to have access to libraries, doctors, and sewage. It is then not surprising to find that these students have lower academic achievement as measured by standardized test scores, persistence and college attendance. This only reinforces the original conditions of inequality.

One factor that generates the difference in teacher quality is that, historically, vacancies in rural schools have been both harder to fill or staffed with relatively lower-quality applicants. In this paper, we show causal evidence that compensation plays a direct role in the observed sorting of teaching talent and the outcomes of students. To establish a causal link between the wages offered at a specific job posting, we leverage a policy change to teacher compensation that raised public-sector teacher salaries by about 25% at 50,000 teaching positions in over 17,000

[1] This is particularly worrying given the recent evidence that teacher quality at all levels has long-term consequences on adult labor market outcomes (Chetty et al. 2014b).

rural schools spread across Perú. The policy, implemented in January 2014, first introduced wage bonuses ranging from S/. 70 to S/.200 to schools located far from the providence capital and with low locality population counts. Arbitrary cutoff rules for policy eligibility generate local quasi-experimental variation in wages across schools. This change in wage structure occurred in the context of a centralized mechanism that assigned teachers to schools based on teacher preferences and teacher performance on standardized competency tests. Specifically, in 2015, the Ministry of Education introduced bi-annual centralized recruitment drives using a centralized system to allocate contract (fixed-term) teachers (*docentes contratados*) and permanent teachers (*docentes nombrado*) across the entire country. This system is unique in that it provides data on job openings and job applications and a known structure regarding how assignments are resolved; these features are not typically observed in most labor markets. These two institutional details present a unique setting in which to study teachers' preferences, and to analyze their sorting patterns across schools/locations with different wage levels.

We start by comparing schools with vacancies in locations around the population threshold of the wage policy. Regression-discontinuity estimates show that teachers who took a position at a rural school with higher wages score 0.7 standard deviations higher in the competency test when compared to teachers who chose a position in lower paying but otherwise similar schools. Teachers in higher paying schools are also more effective – their students perform significantly better in national standardized achievement tests three years after the policy change. We find large and positive impacts on student outcomes at schools that had multiple open vacancies in the previous recruitment drive. In contrast, schools without vacancies experienced small and statistically insignificant effects on student achievement. These two pieces of evidence suggest that it is the inflow of new high-quality teachers that improves student outcomes.

This first set of results suggests that there is no meaningful direct effect of wages on productivity of individual teachers already hired in the system, a finding that is consistent with a recent and related paper studying a large unconditional salary increase in Indonesia by **?**. The authors show that increases in wages have a precise zero effect on student outcomes, and therefore conclude that wage policies are not likely to affect the quality of education. However, in the Indonesian context most teachers are public servants with permanent contracts, thus the selection channel is unlikely to yield relevant effects in the short or medium run. The Peruvian educational system, on the other hand, is similar to the one in other Latin American or African countries, where a large proportion of public sector teaching jobs are staffed by contract, fixed-term teachers. This generates a significant flexibility in the labor market for teachers and a large turnover where the selection margin of wage incentives can play an important role in improving the quality of teachers and student outcomes within a relatively short time pan. As found in other settings (Duflo et al. 2015), the local institutions determining how teachers are evaluated and assigned could be an important necessary condition for increased wages to lead to a meritocratic sorting of talent. The observed policy impact that we measure in Peru may be explained by the pairing of flexible teacher contracts with a transparent, meritocratic assignment mechanism.

These results suggest that targeted pay increases can help reduce spatial inequalities in the access to quality education. However, many factors could contribute to the lower desirability of a location such as the lower levels of school infrastructure and the overall scarcity of services, public goods and local amenities. Therefore, wage policies that adequately compensate for the lack of amenities in rural areas could induce higher-quality teachers to fill positions in less desirable locations. In order to reform teacher pay, policymakers need to know how teachers of different ability levels trade-off pecuniary and non-pecuniary aspects of different job vacancies. If policymakers know the elasticities, they can compute the fiscal costs of counterfactual policies that raise teacher pay in order to raise the quality of teachers in poor rural areas.

To quantify how teachers trade-off compensation and other amenities, we estimate a discrete choice model of the decision of potential teachers to apply for vacant positions in both rural and urban areas. This model allows us to better understand the channels through which the structure of wage incentives shape sorting by quality across space. To the extent that in the assignment mechanism teachers choose their school sequentially based on their ranking in the score distribution, the observed positive effect of wage incentives on the quality of the newly-assigned teachers is consistent with a positive wage elasticity. The estimated model allows us to quantify the magnitude of this wage elasticity and compare it with the other determinants of teachers' demand for job postings, such as their willingness to move/commute away from their current residence and the value of local amenities. We validate our estimated wage elasticity by replicating the observed changes in teachers' scores at the population cutoff that determines eligibility for the wage policy.

We use the model to evaluate the fiscal cost of using wage bonuses to equalize the playing field between children who attend urban and rural schools. Counterfactual experiments suggest that wage bonuses may be an effective policy to make the distribution of teachers more equitable for children attending urban schools versus schools located close to densely populated (urban) areas. However, it is unlikely that wage bonuses would be an efficient instrument to affect sorting at a national scale. Teachers have a strong distaste for moving far away from where they live, which largely outweighs the implied wage elasticity. As a result, it is fiscally expensive to use wage-based policies to equalize the playing field. A less expensive policy option to improve teacher quality in rural areas may be to target teacher training programs and school infrastructure investments in remote, less desirable locations.

These results contribute to the literature on teacher compensation and "pay for performance" schemes (Muralidharan and Sundararaman 2011, Fryer 2013, Barrera-Osorio and Raju 2017, Berlinski and Ramos 2020), showing that relative pay differences can have significant effects on the re-allocation of talent across jobs. In the Peruvian context, teacher compensation is low relative to other college graduates and at baseline it is difficult to staff rural positions with talented teachers. Increasing salaries in this setting is found to generate positive productivity effects through improved ability to recruit relatively more talented teachers. From a policy perspective, this evidence seems particularly appealing to the extent that pay-for-performance reforms are in general less politically viable in the public sector than unconditional wage increases targeted at specific job postings.

More generally, our results are relevant for the design and the evaluation of policies that aim to increase teacher compensation. Several global policy think tanks have recommended for years to increase teacher pay in low-income countries as a way to attract talent towards the education sector (McKinsey 2010, UNICEF 2011, UNESCO 2014). Prior evidence seems to suggest a positive relationship between teacher earnings and school productivity in the long-run (Card and Krueger 1992a,b). However, **?** note that while increasing teacher compensation can improve the overall talent pool through the extensive margin eventually, it may take a long time to see the effects. Furthermore, it will be very costly during the transition if higher earnings do not translate into higher productivity for current teachers as well. This paper addresses a different aspect of teachers' incentive schemes and highlights the notion that not only does the level of compensation matter, but variation within job postings also affects teacher sorting. We show that this channel can have significant effects on the re-allocation of teachers across schools, with crucial implications for the distribution of the quality of education provision.

We also contribute to the recent and rapidly growing literature on the personnel economics of the state (see Finan et al. (2017) for a review). In particular, Dal Bo et al. (2013) show that increased compensation for public sector positions in Mexico lead to a larger pool of applicants, and a higher quality of hired employees. Deserranno (2019) finds that higher financial incentives attract more applicants and increase the probability of filling a vacancy, while crowding out pro-socially motivated agents.

2 Data

In this paper, we use several administrative datasets from the Ministry of Education, which are linked through unique identifiers at either the teaching position level, or at the school level, for each year of our analysis. The data includes administrative panel data on all schools, students and teachers and covers the five year period between 2015 and 2019. We also leverage the centralized system that coordinates the match between all teaching job vacancies and teacher ranked job applications and final assignments. Data on the teacher payroll also allows for the tracking at the microlevel the changes to compensation that arise from recent policy changes. We describe each source of data in more detail in what follows.

The first is the data source stems from the governments *centralized teacher job application and assingment system*. These data include all the vacancies posted at every public school in the country. It also includes teacher job applications to up to five jobs. These are ranked and provided to the centralized system. The data covers the assignment processes that took place in the October 2015 and 2017. These datasets also include information on the teacher evaluations for every applicant in the centralized test, the chosen UGEL, and field of expertise. The data also includes the assignment process for short term contract teachers.

A second administrative data source is the governments *teacher occupation and payroll system (NEXUS).* This is an official dataset used by the Ministry of Education that records all teachers in the Ministry's payroll. It identifies teachers, the school to which she/he is assigned (but not the grade), the type of contract/position (permanent or contract, number of hours, etc.), the base wage and any additional wage bonuses. This information is available for every year between 2012–2018, at the start, middle and end of each school year (March, August and December, respectively.)

A third source of information is the administrative school census dataset that includes *school inputs and characteristics* such as: number of pupils, libraries, computers, classrooms, sport facilities, staff (teachers by status, administrative staff), as well as village-level characteristics: access to basic services (electricity, sewage, water source) and infrastructure (community phone, internet, bank, police, public library). This information is reported yearly by school principals.

A fourth data source is administrative records on *student academic outcomes*. The *Evaluación Censal de Estudiantes* (ECE) is a national standardized test administered at the end of every school year at selected grades by the Ministry of Education to almost all public and private schools throughout the country (coverage is around 98%). We use information on ECE 2016 and 2018 for students in the fourth grade in public primary schools, covering curricular knowledge of math and language (Spanish).[2]

3 Context and Institutions

3.1 Inequality of Education Inputs and Outputs

Perú's colonial history has persistent effects on current institutions, governance, public good provision, and welfare (Dell (2010), Artiles (2020)). The country spans a vast and varied geography made up of mountainous, jungle and coastal regions, with a population composed of a diverse set of ethnic, cultural, linguistic groups who have lived for centuries under extractive

[2]In 2017 there were a large number of floods and landslides throughout the country due to the El Niño natural calamity. This emergency led the Minister of Education at the time to take the (unfortunate) decision to cancel the achievement test for that year.

Figure I: School Infrastructure and Teaching Staff by Rurality

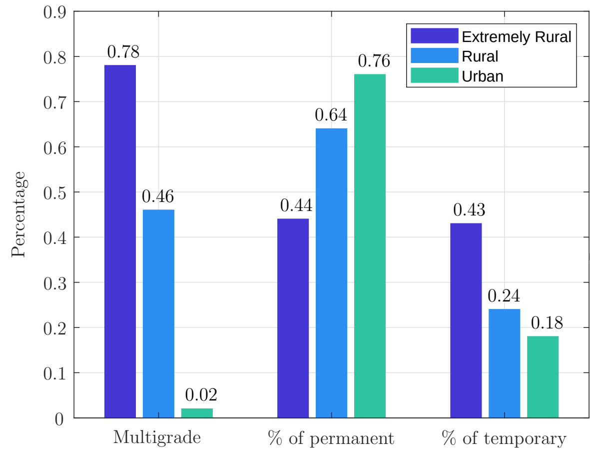

a) School characteristics

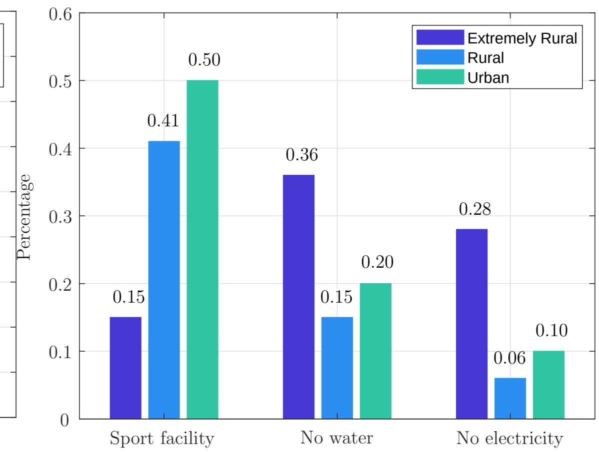

b) Infrastructure

NOTES: These figures show different summary statistics about schools by level of rurality. School infrastructure is shown in the right panel and the type of teachers at the schools in the left panel. *Extreme Rural* schools are shown in purple, other *Moderately Rural* schools in blue, and *Urban* schools in green.

systems of governance as a colony of the Spanish empire (Acemoglu et al. (2001)). The legacy of the colonial institutions and policies is at the root of current structural inequalities: these policies were often targeted the highlands and jungle regions, where most of the natural resources valued at the time were located, and currently coincide with the locations where poverty is concentrated.

Policy makers face staggering differences across urban and more rural communities in their economic development and the access to public services, like education. Overall, Peruvian public primary schools serve 74% of the student population, with the majority of private schools targetting middle class students in urban areas. Instead, in rural areas, public schools are generally the only option to access educational services, and more than 6,000 schools served 98% of school aged children in 2015. Students who attend rural schools score on average 50 points lower on both Math and Spanish exams relative to students in urban schools, are more likely to attend schools serviced by only one teacher ('unidocente'), and attend schools with fewer amenities.

Over the last decade, the government has undergone several efforts to help improve educational attainment in poorer rural areas, such as implementing a large scale conditional cash transfer program, investing in school infrastructure projects, and improving access to drinking water and sewage, among others. However, large differences still exist in the access to education inputs such as school infrastructure and the quality of the teaching staff.

As we describe in the next subsection in detail, in this paper we study the effects of variation in wages across schools in localities with different levels of rurality, as defined by the Ministry of Education (MINEDU) based on the population size and the distance from the province capital.

Figure II: Teacher Competency Evaluation by Geography

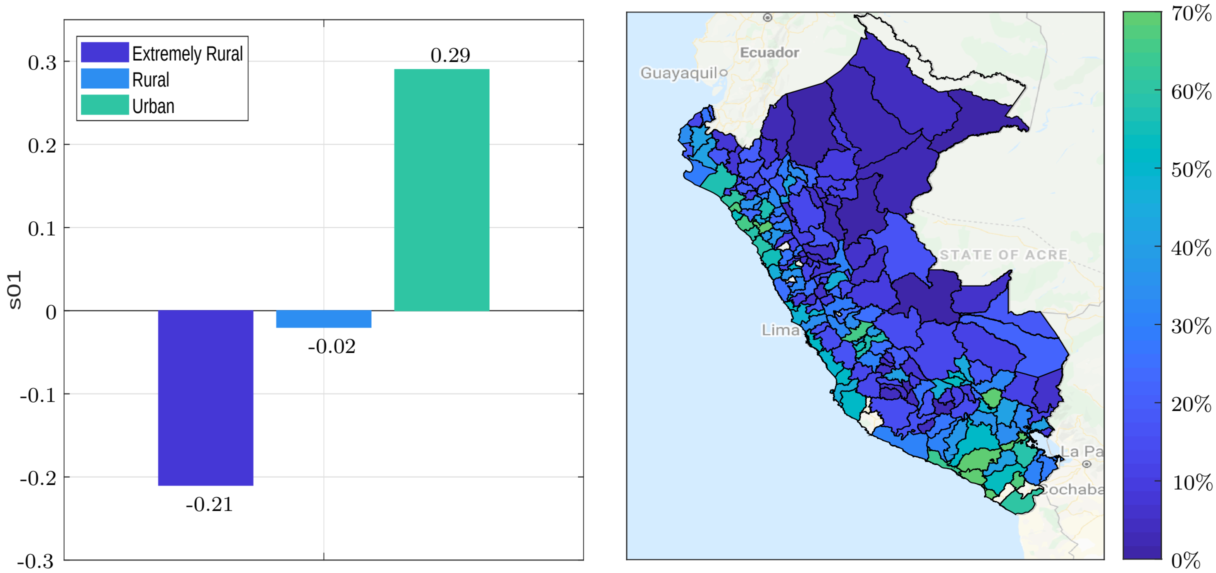

a) Average Teacher Evaluation by Rurality **b) % Teachers with Evaluation Top 25%**

NOTES: This figure shows average teacher evaluatation scores by level of rurality in the left panel. The right panel shows the map of Peru highlighting geographic variation in teacher evaluatation scores. Costal areas are richer and more urban and can be seen to have a larger share of teachers with evaluations in the top quartile.

We now show that there is wide variation in educational inputs across schools by level of rurality. Figure I shows some examples of the stark contrast between schools in rural and urban areas: schools in extremely rural areas are much less likely to have running water, electricity and are unlikely to have any sort of sporting facility. Instruction at rural schools is carried out by teachers who are much less likely to have regular permanent contracts and instead are more likely to have only short term contract teachers or who are less likely to be certified. Table A.1 in the Appendix shows this pattern holds across other indicators for school inputs as well.

Teacher quality has been shown to be an important input into the education production function, and this is true in both developed and developing countries.[3] Teacher subject competency test scores correlated with teacher value added and teacher quality in several contexts (see e.g. Bold et al. (2017), Gregorio et al. (2019), Gallegos et al. (2019).) In our setting, Panel A of Figure II shows that, on average, teachers at rural schools score 0.2σ below the average, while urban teachers score 0.3σ above the average. The geographical spread of teachers with different levels of competency is displayed in the map in Panel B of Figure II, where it is clear that high quality teachers (those at the top 25% of the distribution) are heavily concentrated in the richer, coastal cities, while lower quality teachers are more likely to be in the highlands and the amazon.

[3]Recent evidence has shown this pattern for the US (Chetty et al. 2014a), Ecuador (Araujo et al. 2016), Pakistan (Bau and Das 2020) and Uganda (Buhl-Wiggers et al. 2017).

3.2 Wages, Contracts and Sorting of Public School Teachers in Peru

To better understand the reasons behind the striking inequality in teacher quality across different areas of Peru, we now describe the institutions surrounding the labor market for the nearly 180,000 public school teachers in Perú.

3.2.1 Permanent and Contract Teaching Positions

Public school teachers are hired under two distinct types of contracts. Permanent teachers (*docentes nombrados*) work in conditions similar to tenured teachers in other countries: they are civil servants with permanent contracts, and in practice, the chances of dismissal are close to zero. Teachers can also be hired by the central administration to work at a specific school for the academic year as a contract teacher (*docentes contratados*). These are meant to be one year positions which have the option of renewal for up to one more year. Both permanent and contract teaching positions require teachers to have a diploma (certification), either from a university or technical institute. In cases in which the teaching vacancy not filled through the regular recruitment drive, the school is allowed to hire adults from the local community without teaching certifications.

Contract teaching positions are typically opened whenever local administrations (municipalities) can't secure a long term source of funding to hire a permanent teacher and are conceived as entry-level positions in the teaching career.[4] We will later show that the presence of an opening for a contract teacher is unrelated to the level of wages offered for the opening (or at least that there are no discontinuous jumps with the changes of rurality categories). In Perú, 15% of all primary school teachers were hired as contract teachers, with a much higher share in the most rural schools (43%, see Table A.1).

3.2.2 Public School Teaching Job Compensation

Public school teachers in Perú are paid a fixed wage, and the scale of these wages heavily reward seniority, rather than merit. Teachers' wages depend on (i) the type of contract (permanent or contract teachers), (ii) their seniority, and (iii) the location where they work. In addition to fixed wages, some teachers receive additional bonuses for taking specific responsibilities, for example deputy principals, or for teaching in special education or bilingual schools. In 2015, the average teacher under a temporary contract earned a monthly wage of S/ 1,550 (approximately, US$ 515 using the exchange rate of S/ 3 per USD, from January 2015) in primary school. In contrast, the wages of permanent teachers increase with experience, starting around S/ 1,500

[4] This is a common strategy used in other public educations systems around the world, especially in developing countries. Some notable exceptions are some countries in Southeast Asia, where teachers are included as part of the wider body of civil servants.

and going up to S/ 3,000.

As in most of Latin America, public school teachers' compensations in Perú are low relative to other professionals: the unobserved wage gap between teachers and other professionals with comparable characteristics and educational levels are 30 to 40 percent lower (Mizala and Ñopo 2016). This stands in contrast with institutional settings in other countries in South East Asia, such as India, Pakistan or Indonesia, where public teachers tend to earn more than other comparable professionals (de Ree et al. 2018).

3.2.3 Teaching Vacancies and Teacher Job Applications

Traditionally, the recruitment of permanent and contract teachers in Perú was done in a decentralized fashion. Each year, the government, through its 220 administrative education units (Unidad de Gestión Educativa Local, UGEL), decided the number of positions to be opened. Each UGEL was expected to locally organize the recruitment process following certain guidelines. Little supervision of the process and wide institutional heterogeneity between local administrations generated concerns about a lack of transparency, corruption and political patronage in the hiring of public school teachers. In an effort to make the process more transparent and meritocratic, the Ministry of Education (MINEDU) introduced a nation-wide, centralized recruitment drive, where teacher job postings and teacher job applications were processed on a single, centrally-managed, platform.

The first national recruitment drive took place in October 2015, followed by another round in May 2017. Teachers recruited through the 2015 and 2017 drives started teaching in the 2016 and 2018 academic years (March-December), respectively. At the end of the 2016 (and 2018) academic year, contract teachers had the option to re-apply to their current positions, and their contracts were renewed for an additional year subject to the approval of the school's administration.

The process consists of two main rounds. In the first round, all vacancies for permanent teachers in different fields and specializations are reported to the centralized system. The selection process is designed as an adaptation of a two-sided matching mechanism. Interested applicants must be certified teachers and take a standardized teacher evaluation. This test includes competency on their specific field of expertise, e.g. primary education, secondary math, secondary history and social sciences, etc. Those who passed the minimum required grade were eligible to participate in this first round and could apply for a permanent position. Applicants then choose an UGEL, their field of expertise and, within that subset of job vacancies, they must state their ranked preferences for up to five available positions.

Each school receives a list of up to 20 of the highest scoring teachers among those teachers who ranked an available position within that school in their preferences. The school evaluates the short-list and scores the applicants based on an in-class demonstration, their experience, and an interview. At the end of the process, the grade in the centralized test and the decentralized evaluation are added, and positions are allocated to the highest scoring teacher. There were 19,500 and 37,000 vacancies for permanent teachers in 2015 and 2017, respectively.

In a second phase, all contract teacher job vacancies are provided to the platform, in addition to any permanent positions that remain unfilled in the first process. In contrast to the assignment of permanent teachers, contract teaching positions are designed to be matched quickly by eliminating the schools' (subjective) participation in the screening process. Instead, the mechanism is designed following a serial dictatorship, where schools' preferences are taken to be a strict ranking of the teacher evaluation competency score. Applicants again choose a region (UGEL) and field. They are then sequentially allowed to choose from the available vacancies in that UGEL×field according to their teacher evaluation score. For example, the highest scoring teacher within each UGEL×field gets to choose first among the available teaching vacancies. Once a position is assigned, it is eliminated from the list of available options. The next highest scoring teacher now makes her choice, and so on until either all teaching positions are filled or all teachers are allocated to a position. 56,000 contract teaching teaching positions were available in 2015, and 73,000 in 2017.

Figure III shows the data on applications to vacancies divided by rurality categories. While each vacancy posted at schools in urban areas has many applicants, vacancies posted in rural areas have less than one application, on average (see Panel A). If we further condition on applications from teachers that have an above median teacher evaluation (Panel B), 70% of all extremely rural vacancies have no applicants, while this is true for less than 5% of vacancies posted at urban schools.

Consistent with the application data presented in Figure III, the recruiting process leads to only 35% of job vacancies for permanent positions to be filled in extremely rural schools. At the same time, close to 80% of positions in urban schools were. This is the result of the lack of applications to positions in rural areas, especially when it comes to teachers with higher evaluations and overall qualifications.

As mentioned above, permanent teacher positions that remain unfilled become available for contract teachers (along with other positions initially allocated for that type of contract). Approximately 88% of the positions for contract teachers in rural areas end up being filled by a certified teacher. The remaining 10% of positions are filled in an ad-hoc manner in a decentralized secondary market.

We conclude that the significant inequality in the access to qualified teachers is driven mostly by teacher job application behavior. The microdata on job postings and teacher ranked applications show that most of the applications are concentrated at positions in urban areas, and the system is hard pressed to staff the roughly 17,000 small rural public schools scattered

Figure III: Job Applications by School Rurality

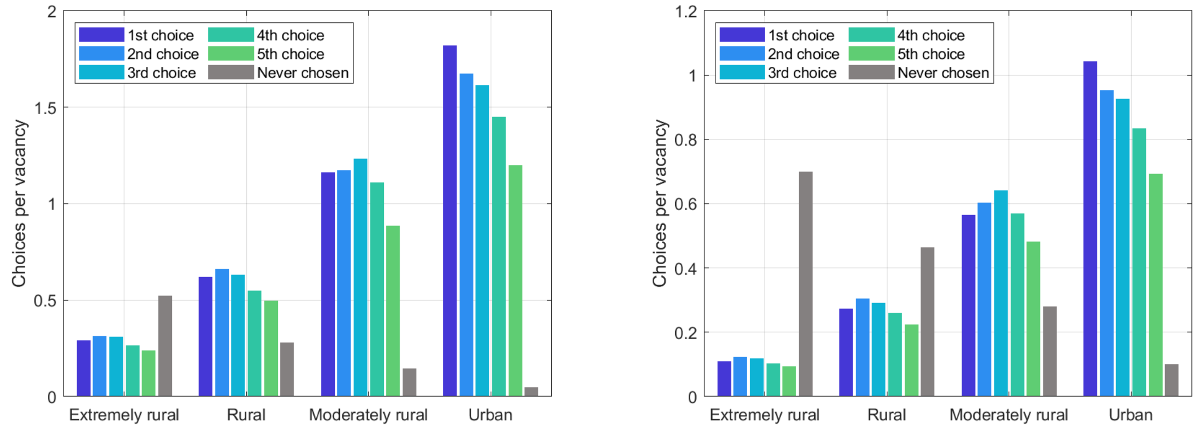

a. Applications for Teachers (Any) *b.* Applications for Teachers (Above Median)

NOTES: These figures show teacher applications for permanent job openings public schools by level of rurality. Application rates for teachers with below median evaluations are shown in the left panel while the right panel shows the application rates for teachers with above median evaluations.

all over the poorest parts of the country. While most of these vacancies are eventually filled using short term contracts, the teachers that eventually take these jobs are those who were not able to find other jobs and were overall significantly less qualified.

3.3 Policy Changes to Compensation in Rural Locations

Teachers in poor rural areas face numerous challenges: scarcity of basic school inputs, lack of services, lack of public goods, few local amenities, and (for some of them) being far from friends and family. Thus the large gap between teachers preferences and the staffing requirements of the state documented above could be due to the fact compensation mostly ignores locality and school amenities associated with the job.

Unlike private labor markets, supply and demand do not determine equilibrium compensation so that if wage setting policies do not adequately compensate for the lack of amenities, those jobs will be less attractive. Consequently, vacancies in rural schools will be either more difficult to fill or will yield less competitive applicants, both outcomes that are consistent with the job application and assignment data presented above.

These considerations motivated the government to introduce a new policy that significantly increased wages at positions in rural schools. Wage bonuses were based on two pre-established criteria that categorized schools into three groups: *Extremely Rural*, *Rural* and *Moderately Rural*. These groups are defined as a function of the population of the local community and the location's proximity to the provincial capital.[5] *Extremely Rural* schools were those located in localities with less than 500 inhabitants, and for which it takes more than 120 minutes to reach the province capital. The second category of *Rural* schools is reserved for either: (a) schools in localities with less than 500 inhabitants and which are located between 30 and 120 minutes from the province capital, or (b) schools in localities with 500-2,000 inhabitants that are farther than 120 minutes from the province capital. The final set of *Moderately Rural* schools are either: (a) schools in localities with 500-2,000 people that are closer than 120 minutes, or (b) schools in localities with less than 500 inhabitants that that are less than 30 minutes away from the capital. All other schools are classified as Urban.

The policy was first implemented in January 2014 providing only *permanent teachers* in *Extremely Rural* , *Rural*, and *Moderately Rural* schools with wage bonuses of S/.200, S/.100, and S/.70, respectively. In August 2015, the bonus for teachers in *Extremely Rural* was increased to S/.500, and wage bonuses were extended to contract teachers. These changes were announced and introduced in August (the middle of the school year) and thus can't affect the selection of teachers prior to the centralized recruitment drive of 2015. The bonus for *Extremely Rural* schools is fairly generous, as it represents 30-40% of the earnings of contract teachers and 20-30% of the earnings of permanent teachers. Figure IV displays the rural categories and the associated wage bonuses as a function of population and time-to-travel as well as the timeline of the implementation of the policy.

Using administrative payroll data for teachers, we can verify the assignment rules lead to significant changes in compensation when schools pass the threshold from one category to another. Figure **??** shows empirically how teacher compensation varies depending on the travel distance from the school to the province capital (left panel), and on the number of inhabitants in the locality (right panel). The Figure is drawn based on teachers payroll data (contract teachers) for December 2015. As can be seen, teacher wages exhibit a large discrete jump when crossing from the left the 120 minutes threshold, and from the right the 500 inhabitants threshold (i.e. the two criteria used to identify *Extremely Rural* schools). The average wage for teachers in schools which are within two hours from the province capital are about S/.300 lower than in schools which are slightly further away. Similarly, the wages drop by S/.200 when the population of the school locality is under 500 inhabitants. When considering both criteria simultaneously (the diamonds in both panels), the observed wage difference between schools missing either of the two criteria and schools meeting both of them is approximately S/. 380. This number matches closely the S./400 (S./430) difference in the wage bonus between *Extremely Rural* and *Rural* schools.

[5]The population of the locality where the school is located is measured by population counts in the latest available census. The time it takes to travel from the locality to the province capital is measured on the basis of GPS coordinates taken by an inspector from MINEDU, after taking into account usual modes of transport and types of roads available at the time of the measurement.

Figure IV: Spatial Distribution of Rural Schools

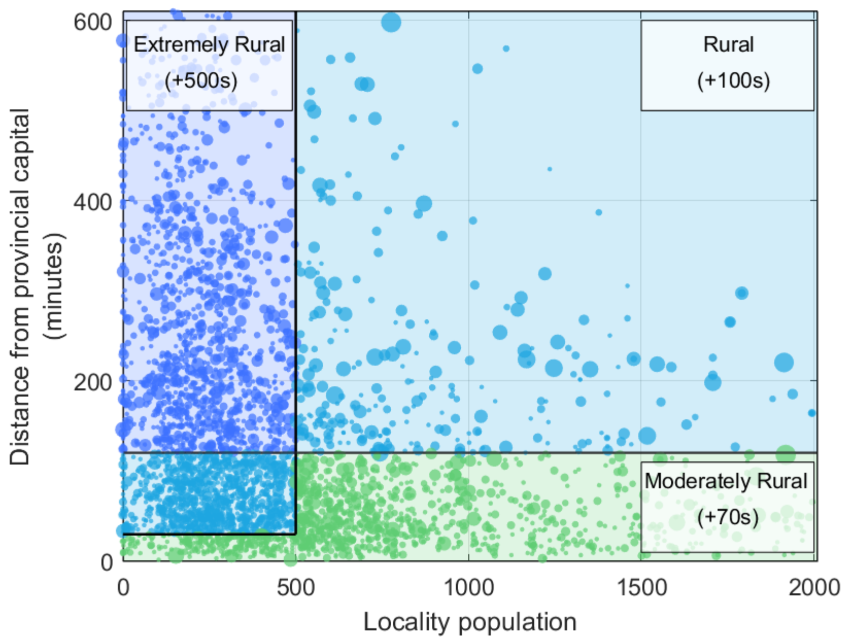

Notes. This figure shows the spatial distribution of rural primary schools along the two dimensions that determine the assignment of the wage bonus. *Extremely Rural* schools are the purple dots, *Rural* are light blue and *Moderately Rural* schools are green. The size of the dots reflects the cumulative number of open vacancies in each school over the recruitment drive.

In what follows, we provide descriptive data on the potential dimensions in which the wage policy could have affected the selection of new recruits in public sector schools. To do this, we use data from the contract teacher recruitment drives from 2015 and 2017, and show that there is a clear association between the level of wages and (i) the probability that a vacancy is filled, (ii) a proxy for the demand for vacancies, and (iii) the quality of teachers choosing a certain school. Figure V shows the coefficients (and 95% confidence intervals) from an OLS regression in which we use as a dependent variable different measures of the three dimensions above, on dummies for whether the school is in an Urban area (no bonus), in a rural area with a small (S/ 100) bonus, or a rural area with a large bonus (S/ 500), controlling for UGEL and year fixed effects (and cluster standard errors at the school$\times$year level).

Panel A in Figure V shows that schools that pay higher wages are also those that are more likely to fill up a position advertised for contract teachers. While schools with no bonus or low bonuses fill about 88% of positions, paying a high bonus is associated with about 7% more vacancies filled. This is prima facie evidence supporting the hypothesis that higher wages may have an effect on the demand for teaching positions, which in turn, can lead to a reduction in the structural inequalities in education. In Panel B, we test for the association with a more direct measure of the demand for teaching positions. Here we use the relative rank (normalized between 0 and 1) in which a position has been filled through the differed acceptance mechanism.

Figure V: Contract teacher selection

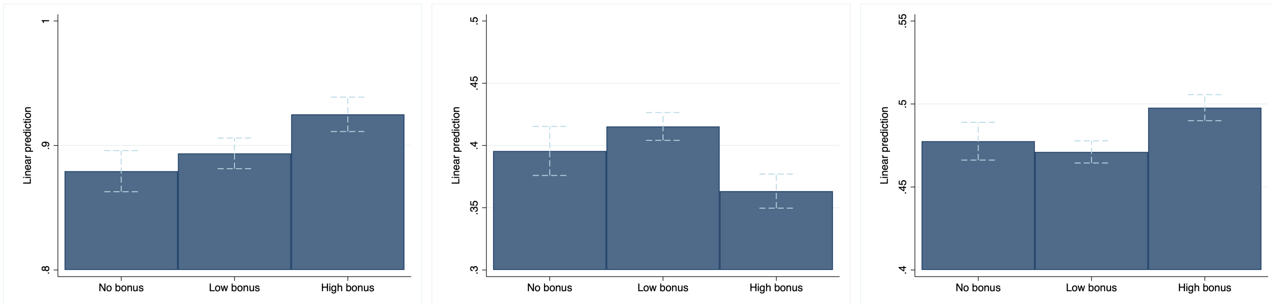

a. Vacancy filled

b. Teacher rank

c. Teacher score

Recall that the contract teacher recruitment is a tournament between applicants within an UGEL and specific field of expertise. We compute the the rank in which a position is filled by counting the positions in the ranking in an UGEL×field, and dividing it by the total number of people in the tournament. The results show that high paying positions tend to be filled much faster, compared to those that offer either a low bonus or no bonus at all.

Next, we examine whether the increase in the demand for teaching positions translates into higher quality teachers being attracted to teach at schools that offer a higher compensation. We do this by using the results from the centralized tests as part of the recruitment process. Again, here we see a similar pattern, where high bonuses are associated with test scores that are much larger than those obtained by people who choose a school with a low or no bonus. We revisit these associations and establish causality on these results in the next section.

Finally, we expect not only the level of the bonus to matter for the demand for positions, but also the size relative to the base wage. This is relevant because from here we would be able to derive relative elasticities that then we will map to our structural estimations in Section 5. In Appendix Figures A.2, A.3 and A.4, we show the same patterns as before, but splitting the sample by the base wage of each teacher. The results show that the same patterns as before hold, but there is a clear association between the size of the effects and the magnitude of the increase in wages due to the bonuses.

4 Policy Effects of the Increase in Compensation

4.1 Identification and Estimation

Offering higher wages for teachers who choose to take a position at a rural location could attract more and higher-quality teachers, in turn leading to a reduction in educational opportunities inequality. On top of the potential selection effects induced by higher wages, additional compensation could increase the quality of instruction by affecting the productivity of the teachers who are already employed in rural schools. Alternatively, higher wages may not change teachers' application behavior if they are not enough to compensate for the local amenities or distance from their current location.

In sum, for a wage-reform policy to improve access to educational opportunities in poorer rural areas, three conditions must hold: first, compensation must have an effect on teachers' job application behavior so as to generate changes in teachers' sorting across space. Second, the increase in demand for school vacancies must lead to an increase in teacher quality through the assignment system. Third, the increase in teacher quality must have an effect on student achievement.

In the empirical analysis that we discuss in this section, we will trace out each of these conditions. We exploit the discrete changes in the assignment rules of the wage bonus policy in order to identify the causal effects of unconditional wage increases on (i) the demand for teaching positions across the entire country, (ii) the selection of the assigned teachers into the job postings, and (iii) students' standardized test scores. Our main estimating equation is as follows:

$$y_{jt} = \beta_0 + \beta_1 ExtRural_{jt} + f(pop_{jt}, time_{jt}) + \delta_t + \epsilon_{jt}, \tag{1}$$

where y_{jt} is an outcome variable for school j at time t. The treatment is defined by *ExtRural*, an indicator variable that is equal to one if school j's locality has less than 500 inhabitants ($pop_{jt} < pop_c$) *and* it is located more than 120 minutes away from the province capital ($time_{jt} > time_c$). In our main specification, we control for flexible polynomials $f(\cdot)$ of the running variables. The parameter of interest is β_1, which captures the effect of wage bonuses on teacher or student outcomes. We pool data from the two centralized recruitment drives (and the subsequent school years), therefore δ_t is a time dummy indicating the specific year of the recruitment drive (for teachers) or the school year (for students). The error term ϵ_{jt} is clustered at the level of assignment of the wage bonus: the school-year pair.

The validity of this sharp RD design relies on two testable implications. First, the density of observations has to be continuous around the threshold, implying that there was no endogenous sorting of schools around the threshold. Second, schools around the threshold are similar in all observable characteristics at baseline.

The policy under study may have generated incentives for school principals and administrators to partly manipulate some of the information required for the assignment rule, thereby leading to a violation of the first assumption. We test this empirically in Figure B.1 of the Appendix, where we display the densities based on local-quadratic density estimators with the corresponding confidence intervals for each of the assignment variables in each of the two years of the assignment mechanism. The population threshold is based on census data, and as such it is difficult to manipulate. The left-hand side panels show that indeed, there are no significant discontinuities at the 500-inhabitants threshold for either of the years of interest. The panels at the right hand side of Figure B.1 shows the empirical densities of observations around the time-to-travel distance threshold. Instead, there is a significantly larger mass of schools that fall just above the time-to-travel threshold for the assignment mechanism that took place in 2017. The formal manipulation tests (Cattaneo et al. 2020) confirm these visual patterns.[6] Time-to-travel information is gathered by inspectors from the Ministry of Education, who physically go to the schools and take a GPS measurement of the school's location. The GPS measurement was updated in 2017, and by that point, the previous measurement was public information, which provides large incentives for schools close to the threshold to manipulate the measurement and gain access to the wage bonus for all of their teachers.

Overall, the data shows that schools may be sorting endogenously across the time-to-travel distance threshold, whereas there seems to be no strategic manipulation of the population assignment variable. This pattern is also evident when we do the validity checks for the second assumption, namely, whether there is balance in observables across the thresholds. Table B.1 shows that school and locality-level covariates are smooth around the population threshold. Columns (1) and (4) report RD estimates of the empirical specification in Equation (1) for the population discontinuity, separately for schools that had an open position in the 2015 and 2017 assignment process, respectively. The point estimates for the β_1 coefficient are very small and not statistically different from zero in all but two cases for 2015, and in all but three cases for 2017 (out of the 32 covariates considered). We get a similar result when limiting the sample to schools that had an open position for permanent teachers or for contract teachers during the two years of the centralized recruitment drive, which are displayed in columns (2),(3),(5), and (6). Given the manipulation of the time-to-travel threshold, we only rely on the source of variation in teacher wages provided by the population threshold for our main estimation.[7] More precisely, given continuity of potential outcomes around the population cutoff, the following variant of

[6]The estimated bias-corrected t-statistic for the null hypothesis of no difference in height between the two density estimators for the time-to-travel discontinuity is 2.33 (p-value=0.02) in 2017 and 1.53 (p-value=0.13) in 2015. T-stats are lower in size and they are not statistically significant for the population discontinuities, taking the value of -1.07 (p-value=0.28) in 2015 and of -0.12 (p-value=0.90) in 2017.

[7]The fact that the time-to-travel variable is manipulated in 2017 does not prevent us from using the (unmanipulated) 2015 time-to-travel as the running variable for 2017. However, note that there were many changes in the assignment of the rurality categories between 2015 and 2017. Figure B.2 in the Appendix illustrates this point by plotting the assignment to the *Extremely Rural* category in each year agains the running variable of the opposite year. Unsurprisingly, the population category does not change between years (given that it is based on census data), and the running variable of either years is a good predictor of the treatment status. In contrast, time-to-travel in 2015 does not help predict the treatment status in 2017, and therefore doesn't provide useful variation to estimate the effects of the wage bonus in 2017.

equation (1) identifies the effect of the wage bonuses:[8]

$$y_{jt} = \gamma_0 + \gamma_1 \mathbf{1}(pop_{jt} < pop_c) + g(pop_c - pop_{jt}) + \delta_t + u_{jt}, \tag{2}$$

where, as before, $g(\cdot)$ is a flexible polynomial of the distance from the population cutoff and u_{jt} is an error term clustered at the school-year level. We estimate γ_1 non-parametrically using the robust estimator proposed by Calonico et al. (2014) through local-linear regressions that are defined within the mean-square error optimal bandwidths. Our main interest is on the γ_1 parameter, which represents the average difference in the outcome of interest between schools, teachers, or students in localities that are just above or below the population threshold, and therefore that are marginally eligible to receive (or not) the wage bonus. The estimator identifies the effect of crossing the population threshold at different points (see Figure IV), and therefore represents a weighted average of (i) the effects of a S/ 400 bonus implied by comparing Rural and Extremely Rural localities, (ii) the effect of a S/ 30 bonus due to the comparison between Moderately Rural and Rural localities, (iii) the effect of a S/ 430 bonus, implied by the comparison of Moderately Rural and Extremely Rural localities.[9]

4.2 Compensation and the Demand for Teaching Positions

As shown in Figure **??**, the average contract teacher in a school located in an *Extremely Rural* place earns S/ 380 more than those in a school in a locality that is just above the 500 inhabitants threshold, which on average represents an increase in the unconditional wage of 23%. This wage increase makes teaching positions more desirable, increasing the demand and potentially leading to a better selection of personnel (Deserranno 2019, Dal Bo et al. 2013). We analyze these margins of response in Table I. The corresponding graphical evidence is shown in Figures B.3 and B.4 in the Appendix.

We start our empirical analysis by studying whether wage bonuses affect teachers' preferences for choosing in which school to work, which is a precondition to observing effects on teacher location choices and schooling outcomes. To study the effects on the demand for teaching positions, we make use of the revealed preference data obtained from the applications for permanent teaching positions. Recall that to be selected for these positions, applicants go

[8]Note that we exclude from the estimation localities within 30 minutes of the province capital, since for them, crossing the population threshold does not imply an increase in the bonus.

[9]An alternative approach to the regression model depicted equation (2) would be to limit the sample to schools located above the 120 minute time-to-travel threshold, and compare only Rural and Extremely Rural localities. We don't pursue this empirical strategy for two important reasons. First, restricting the sample to schools located above the time-to-travel threshold would imply conditioning on a variable that is partially manipulated (see Figure B.1). Second, such sample restriction would also imply leaving out a large portion of schools, and in particular those that cross the threshold along the diagonal of Figure IV, i.e. from *Moderately Rural* to *Extremely Rural*, thereby missing relevant variation of wage bonuses in the data.

through a two-sided matching mechanism in which those who scored above a certain threshold in the centralized test had to submit their preferences for positions, then the top candidates who preferred a certain school were evaluated on-site. Ultimately, the grades from both evaluations determined who was offered the job.

In Columns 1 and 2 in Panel A of Table I we study the effects of the wage bonus on teacher's preferences. Here, we follow the empirical specification in Equation 2, using a dataset at the school level, and the dependent variable is either a dummy that takes a value of one if at least one teacher in our sample included the school in their ranked-order lists (column 1), and the maximum preference-rank of the school among all teachers in our sample (column 2). We find that the wage bonus has a large and positive effect on the demand for teaching positions. In our sample, the average school is mentioned in 76% of teachers' lists and this proportion increases by 23 percent for schools just below the population threshold. We do not see a similar effect when looking at the maximum preference-rank of the school, although the interpretation of this result is complicated by the fact that this regression is estimated on the (selected) sample of schools that were mentioned by at least a teacher.[10]

The two-sided matching mechanism implemented does not ensure that the preferences rankings submitted are strategy proof. Indeed, teachers may submit ranked-order lists (RoLs) that are strategically chosen to maximize their chances to get their preferred position. However, if it is indeed the case that teachers believe that positions that offer higher wages are more competitive, they should mention them in their list *less* often, making our estimates a lower bound for the demand for positions that offer higher levels of compensation. Overall, the fact that we see that teaching positions entitled to a higher wage bonus are mentioned in teacher's list more often than those that receive a lower bonus is a strong indicator that higher compensation increases the demand for these teaching positions.

4.3 Compensation and Recruitment of Talent

More demand, and therefore more competition, for positions can lead to an increase in the quality of applicants who select into higher paying teaching jobs. This effect can potentially be explained by two different mechanisms. On the one hand, higher wages generate a larger demand for positions (as shown in subsection 4.2) and attract a larger pool of applicants, increasing the average quality of the marginal applicant. Alternatively, an increase in the quality of the marginal candidate taking a position at a high-paying school could be explained by pure sorting *within* the system, whereby higher ability teachers who would have otherwise

[10]Schools that receive a wage bonus are much more likely to receive a preference (see Column 1), leading to differential sample selection at the two sides of the threshold. To deal with this issue, Table I reports the bounds of the effect estimated following (Gerard et al. 2020). The estimated bound suggests that higher paying schools are mentioned up to 0.6 positions earlier, but its upper limit (0.3) exceeds zero.

Table I: Monetary Incentives and Teacher Selection

Panel A: Permanent teacher				
	(1)	(2)	(3)	(4)
	In ROL	Highest preference	Vacancy filled	Teacher score (std)
Above cutoff	0.177***	0.004	-0.001	-0.014
	(0.068)	(0.234)	(0.071)	(0.175)
Bounds	[.154; .28]	[-.61; .363]	[-.005; .041]	[-.271; .201]
Mean dep. var. (LHS)	0.759	1.838	0.372	-0.099
BW	173.292	148.865	165.291	205.787
Schools	860	733	830	1052
Observations	1038	700	1697	864
Panel B: Contract teacher				
	(1)	(2)	(3)	(4)
	Vacancy filled	Teacher rank	Teacher score (std)	Top 25%
Above cutoff	0.043	-0.120***	0.451***	0.127***
	(0.045)	(0.035)	(0.122)	(0.049)
Bounds	[.047; .047]	[-.13; -.11]	[.4; .468]	[.126; .126]
Mean dep. var. (LHS)	0.898	0.372	0.068	0.225
BW	160.098	158.021	162.492	200.874
Schools	943	930	961	1203
Observations	2218	2020	2083	2793

NOTES. This table reports the effect of crossing the population threshold on different outcomes. Panel A uses the sample of permanent teachers. In column (1) the outcome variable is a dummy equal to one if a school was mentioned in at least one application, while column (2) uses the highest preference expressed in a teacher's list. Column (3) studies whether a vacancy was filled by a certified teacher, and Column (4) uses as outcome variable the standardized test score obtained in the centralized test. In Column (2) the sample is restricted to schools that were mentioned in at least one application, while in Column (4) is restricted to vacancies that were actually filled by a certified teacher. Panel B focuses on the selection process of contract teachers. Columns (1) and (3) are analogous to Columns (3) and (4) from panel A. Column (2) shows the effects on the rank in which a vacancy was chosen in the deferred acceptance mechanism (normalized so that it takes value from zero to one). Column (4) is an indicator equal to one if the teacher selected in a certain vacancy scored in the top 25% of the distribution and zero for bottom 75% and non-certified teachers. Cells report the bias-corrected regression-discontinuity estimates obtained using the robust estimator proposed in Calonico et al. (2014) and their bounds estimated using the procedure developed in (Gerard et al. 2020). Regressions are defined within a mean-square error optimal bandwidth (BW), reported at the bottom part of the table. The table also reports the mean of the dependent variable computed within the intervals $(0, +BW)$ (right-hand-side of the cutoff) and $(-BW, 0]$ (left-hand-side of the cutoff). SE are clustered at the school×year level. *** p< 0.01, ** p<0.05, and *p<0.10.

gone to *Moderately Rural* (or even Urban) schools are instead choosing *Rural* or *Extremely Rural* positions due to the wage incentives. In this subsection, we explore these two hypotheses by looking at permanent and contract teachers.[11]

A first-order dimension of the effect of higher wages is whether or not a position is filled or not, either by a permanent or contract teacher. Having a teacher recruited through the merito-

[11]We present all the results pooling the data from the two recruitment drives from 2015 and 2017 in the main text. The results split by year are shown in Tables B.3 and B.4 in the Appendix. Overall, the broad patterns described here hold for both rounds of the recruitment drive. Similarly, Tables B.5 and B.6 in the Appendix break up the results estimating the effects for 'high' and 'low bonus', i.e., going from *Moderately Rural* or *Rural* to *Extremely Rural* or from *Moderately Rural* to *Rural*, respectively. As expected, most of the effects discussed in this section are concentrated in the 'high bonus' category.

cratic selection process mechanically increases the quality of the selected teacher. Moreover, the worst case scenario in which the vacancy goes unfilled would imply an increase in the workload for other teachers in the school, presumably, reducing their effectiveness. Column (3) of Panel A in Table I presents the results for the probability that a vacancy is filled in the permanent teacher selection process, while column (1) of Panel B shows the analog for contract teachers. For this analysis, we use the sample of all open vacancies for each type of position. While higher wages increase the demand for permanent teaching positions (see subsection 4.2), the effect of wage bonuses on the probability of a filled vacancy is a precisely estimated zero. We also see a precisely estimated zero effect when studying the effect of the wage bonus on the quality of permanent teachers who take a position at a high-wage school (as measured by the score in the standardized test), as shown in column (4) in Panel A of Table I. While surprising at first, this null effects are likely due to the fact that the decentralized stage of the selection process – when schools have discretion to assign grades based on the on-site interviews with the candidates – completely undoes the effects of the meritocratic selection of the centralized stage. Supply-side responses may thus have neutralized the demand-side effects triggered by increased wages. This result underscores the relevance of transparent, objective and meritocratic selection practices for human resources in the public sector (Deserranno et al. 2021).

For contract teachers, who are selected through the serial-dictatorship mechanism described in Section 3.2, column (1) in Panel B of Table I documents a positive but statistically insignificant effect of higher wages on the probability that a teacher accepts a position.[12] The nature of the assignment mechanism prevents us from directly observing the demand for teaching positions of contract teachers, as it was the case for permanent teachers. However, we can indirectly measure the demand for higher-paying short-term teaching positions by comparing the rank order in which they were filled. In column (2) in Panel B of Table I we measure the rank order, normalized to be between zero (filled by the first applicant in the UGEL $\times$ Field) and one (filled by the last applicant in the UGEL $\times$ Field). On average, a standard teaching position is filled by a teacher ranked in the 37 percentile, while schools that offer a wage bonus manage to fill the position with a candidate in the top quarter of the distribution of applicants. Consistent with the evidence from the previous subsection, this result suggests that higher compensation

[12]The estimation sample consists of all the open positions in both recruitment drives for contract teachers. This sample is the union between positions that didn't get filled in the permanent teacher selection process and those that were only opened for contract teachers. To ensure that we do not estimate our treatment effects on a selected sample, Figure B.5 in the Appendix shows that the density of open vacancies does not discontinuously jumps at the threshold.

increases the demand for contract teacher vacancies.[13]

In columns (3) and (4) of Panel B in Table I, we evaluate whether the increased demand for high-paying positions leads to an increase in contract teacher quality. Teachers who select into schools that offer a higher wage bonus have a score in the evaluation test that is 0.45 standard deviations higher than those who choose a position in another rural school. This is not a marginal increase in quality, as we show in the column (4): newly recruited teachers are 12.7 percentage points more likely to be in the top 25% of the quality distribution. This finding is presented in a more general fashion in Panel B of Figure VI, where we show the estimates on the probability that a teacher scoring in each decile of the quality distribution takes a position at a high-paying school. The results show that higher compensation not only makes it more likely that a high scoring teacher (i.e., in the 7th to 9th deciles) chooses a high paying school, but this also makes it less likely that a low scoring teacher (i.e., in the lower half of the distribution) takes a job in these schools.[14] This sharp change in the quality distribution of new teachers explains the large point estimate observed in column (3) of Panel B in Table I. In sum, higher compensation causes an increased demand for contract teaching positions, and these effects are accompanied by a large and significant increase in the selection of high-quality teachers into these schools.[15]

As mentioned earlier, the finding that the marginal teacher who selects a high-paying position is of higher quality could be due to the increase in demand and higher average quality or applicants, or due to sorting within the system. Having teachers who take positions just below the population threshold being those who would have otherwise chosen a school just above the threshold is problematic as it would imply a violation of SUTVA (Rubin 1986) in the context of our RD design. We address this issue in Figure VII (and its companion Table B.9 in the Appendix). We run the regression model depicted in equation (2) using as dependent variable an indicator which takes the value of one if a teacher who accepts a position at a school just

[13] Note that the results in columns (2)-(3) in Panel B of Table I are estimated on the sample of schools for which a vacancy was filled. While for permanent teachers this does not cause significant concerns, for contract teachers the probability of filling a vacancy slightly increases for high-paying schools (although not significantly – see column (1) in Panel B of Table I). To deal with this concern, we bound our estimates using the method outlined in (Gerard et al. 2020), which we show below the point estimates of Table I. In all cases, the bounds are quite tight and all of our conclusions hold. For completeness, we also present the bounds for all other regressions in the table.

[14] For completeness, we show a similar analysis for permanent teachers in Panel A of Figure VI, where, as expected, we see a zero effect throughout the teacher quality distribution.

[15] In Table B.8 we study whether higher wages systematically attract contract teachers with specific characteristics. Our results suggest that teachers who select into higher paying positions are more likely to be female (column 1, but imprecisely estimated), about 1.5 years younger (column 2), and 5.2 percentage points more likely to be novice teachers in the public school system (column 3, imprecisely estimated). This is consistent with the fact that the probability that these teachers have more than 3 years of experience drops by about 14.5 percentage points (column 5). Taken together, these different pieces of evidence suggest that the vacancies in higher paying positions are partly being filled by new comers that are drawn into the public education system by the wage incentives rather than a pure reallocation effect of the policy within existing teachers.

Figure VI: Wage Bonuses and the Selection of Quality Teachers

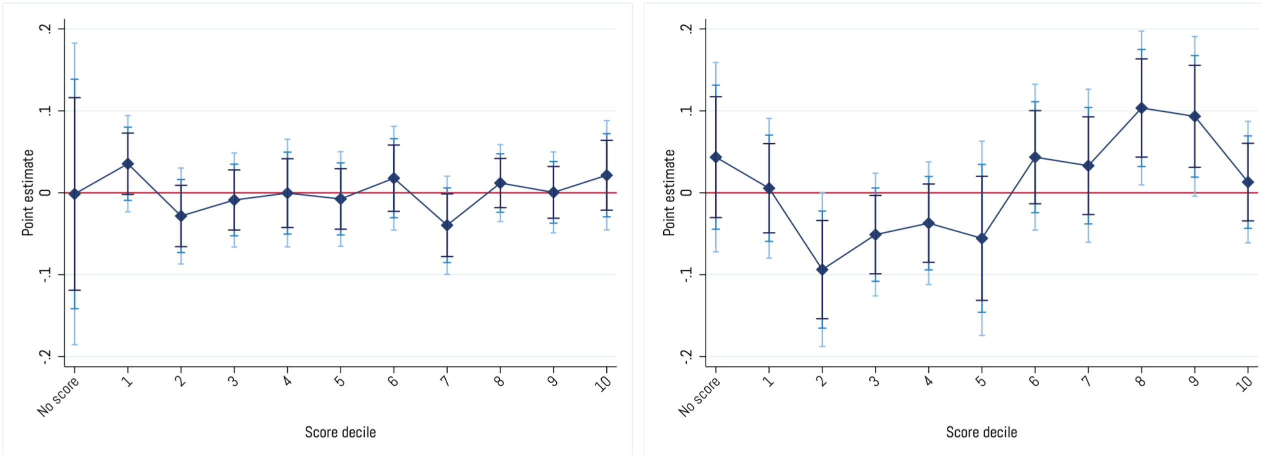

a. Permanent teacher *b.* Contract teacher

Notes. The figure displays the effect of crossing the population threshold on different measures for the quality of recruited teachers. These are a dummy equal to one if the vacancy is filled by a non-qualified teacher, and a set of binary indicators for whether the vacancy is filled by a teacher whose score falls into the decile of the score distribution reported on the x-axis. In panel (a), the sample includes only schools with a permanent teacher opening in the 2015 or 2017 assignment process, and the score distribution is defined based on the sample of teachers who participated in the assignment process for permanent teachers. panel (b) is the analogous of (a) for the selection process of contract teachers. Non qualified teachers are defined as teachers who did not pass the minimum required grade for a permanent position (panel a), and teachers without a score in the standardized test (panel b). Markers and vertical lines indicate the robust bias-corrected regression-discontinuity estimates and confidence interval (at the 95% level) obtained using the robust estimator proposed in Calonico et al. (2014). Regressions are defined within a mean-square error optimal bandwidth.

below the population threshold was previously teaching at a school in a location within a certain population range or whether or not she/he was a new comer to the public school system. Most coefficients reported in Panel A of Figure VII are close to zero, and the large majority of them are not statistically significant. This evidence suggests that the observed increase in teacher quality is not the result of a zero-sum game among schools located across the population cutoff, but rather that the wage bonus attracts a quite diverse pool of applicants either coming from schools in localities of a wide range of population size (including urban areas) or new entrants in the public education system. On the other hand, Panel B deals directly with the concern that our empirical design violates SUTVA. We plot the predicted probability that a teacher coming from a certain population bin ends up taking a position at a school just below the population threshold. The data shows that overwhelmingly, new comers to the profession and teachers already in the same school are more likely to take these positions (light blue bars), compared to those coming from schools that are not allocated a wage bonus (dark blue bars).

4.4 Wage Bonuses and Student Achievement

In this last subsection, we study the extent to which the improved selection of contract teachers due to higher wages leads to improvements in students' academic achievement. While the

Figure VII: Origin of Newly Recruited Teachers

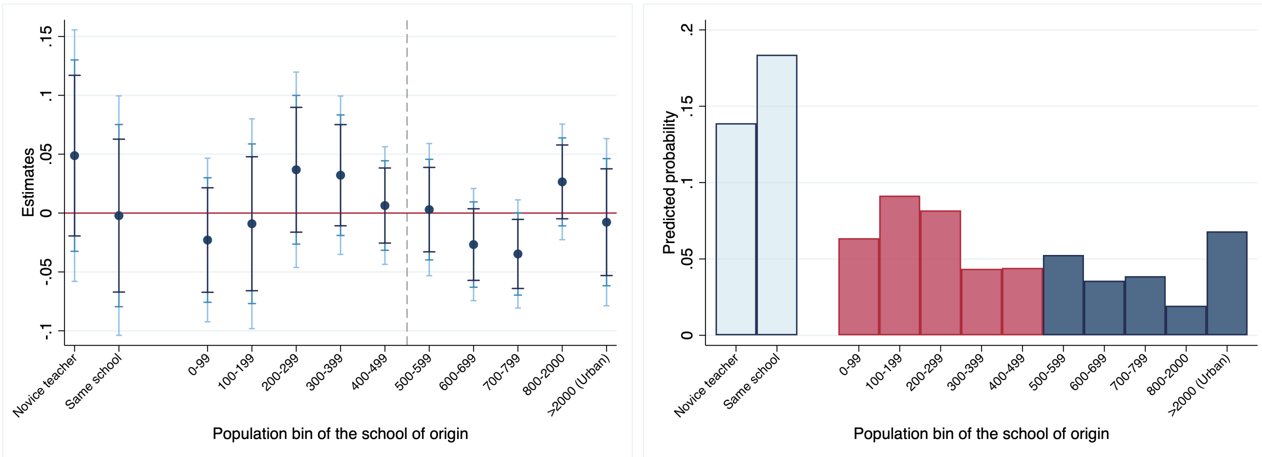

a. Regression Estimates *b.* Levels

Notes. This figure displays the effect of crossing the population threshold on a set of indicators for the teachers' location in the year before the assignment process. These are a dummy equal to one if the vacancy is filled by a new entrant in the public education system, is filled by a teacher already in the same school, or is filled by a teacher whose previous location falls into the population bin indicated in the x-axis. Urban schools are those in localities above 2000 inhabitants, Teachers' previous school (and the status of new entrant) is determined based on the teacher occupation and payroll system (*NEXUS*). The sample includes only schools with a contract teacher opening filled in the 2015 or 2017 assignment process. Markers and vertical lines in panel (a) indicate the robust bias-corrected regression-discontinuity estimates and confidence interval (at the 95% level) obtained using the robust estimator proposed in Calonico et al. (2014) (also reported in table format in Appendix Table B.9). Vertical bars in panel (b) indicates the estimated left-hand-side intercept, that is, the estimated probability that a vacancy is filled by a teacher whose previous school falls into the corresponding category, in vacancies just below the population threshold.

average school in our sample has few teachers (about 3), the data available does not allow us to precisely match teachers with a specific class within a school, and hence we are unable to pin down the direct effect of having a better teacher (due to higher wages) in the classroom. Instead, we show the 'total policy effect' by comparing students' achievement in schools where teachers received a higher wage bonus vs. those that were not eligible.

Recall that wage bonuses to teachers in rural schools are not restricted to those who are recruited through the centralized recruitment drive, but rather they affect all teachers in the school. Hence, as mentioned in subsection 4.1, these bonuses could potentially affect student achievement through two main mechanisms: (i) increased teacher effort due to a higher compensation, or (ii) improvements in the quality of newly selected teachers.[16] We explore two these mechanism in Table II, where we use student level data on the results in the Math and Spanish tests in 2018 for fourth graders. Test scores are standardized, and we consider here

[16] In principle, wage bonuses could also affect student achievement by changing the size of the teaching staff. However, in Appendix Table B.7 we show that the wage policy has a small and statistically insignificant effects on teachers' number and students to teachers ratio.

Table II: Wage Bonus and Student Achievement

	No vacancy	Vacancy		
	(1)	(2)	(3)	(4)
		Any vacancy	Permanent teacher	Contract teacher
Above cutoff	0.028	0.346**	-0.058	0.421***
	(0.172)	(0.141)	(0.213)	(0.155)
Mean dep. var. (LHS)	-0.477	-0.466	-0.339	-0.485
BW	124.728	109.312	202.161	106.309
Schools	372	696	349	588
Observations	3948	9771	3979	8464

Notes. This table reports the effect of crossing the population threshold on student achievement. In all columns, the outcome variable is the average of the standardized 2018 test scores in Math and Spanish for students in fourth grade. The sample in Columns and (2) is split based on whether the school had an open vacancy (of any type) in the 2015 or 2017 centralized recruitment drive. In Column (3) and (4), the sample is further restricted to schools that had vacancies for permanent or contract teachers, respectively. Each cell reports the bias-corrected regression-discontinuity estimates obtained using the robust estimator proposed in Calonico et al. (2014). Regressions are defined within a mean-square error optimal bandwidth (BW), reported at the bottom part of the table. The table also reports the mean of the dependent variable computed within the intervals $(0, +BW)$ (right-hand-side of the cutoff) and $(-BW, 0]$ (left-hand-side of the cutoff). SE are clustered at the school×year level. *** p< 0.01, ** p<0.05, and *p<0.10.

the average score in the two tests as the main outcome of interest at the student-level.[17]. If wage increases cause an increase in teachers' effort, which in turn lead to improvement in student achievement, we should observe that student outcomes also improve in schools that didn't have an open vacancy. Column (1) in Table II explores this hypothesis by looking at students' achievement in 2018 in the sample of schools that did not have an open vacancy in the 2015 or 2017 recruitment drives. For these schools, the wage bonus does not significantly affect student achievement: the estimated RD effects are very small and statistically insignificant. This result is consistent with the findings in de Ree et al. (2018), where they show that doubling the wages received by teachers who were already working in Indonesian schools at the time of the reform didn't cause any improvements in students' outcomes (but increased teachers' life satisfaction outcomes).

In column (2), we restrict the sample to schools that had an open vacancy in 2015 or 2017, either for a contract or permanent teacher, while in columns (3) and (4) we separate the sample

[17]The results for each subject are also show in Table B.10 in the Appendix. By 2018, students in these schools have been exposed to teachers recruited in 2015 for two years and to teachers recruited in 2017 for one year. Given that we are not able to match students to their teachers, this increases the chances that the students we are able to observe in the standardized tests are exposed to a teacher recruited through the centralized recruitment drive. The centralized test in 2017 was cancelled due to climatic conditions (*El Niño* phenomena hit large areas of the country, causing classes to be cancelled in many regions).

between schools that had vacancies for permanent or contract teachers, respectively.[18] Overall, students in schools that had any open vacancy performed much better if teachers received a higher wage: the effect size is of 0.3 standard deviations. Consistently with the null selection effect of wage bonuses on permanent teachers reported in Panel A of Table I, in column (3) we show that the effect of the wage bonus on student performance is very small in magnitudes and statistically insignificant.

In column (4) of Table II we focus on those schools with an open vacancy for a contract teacher. Consistent with the substantial increase in teacher quality as a result of the wage bonus policy, students in schools in which teachers were paid a wage bonus, students score 0.4 standard deviations higher in standardized achievement tests, compared to those in schools that also had an open vacancy, but were not eligible to receive a bonus. The magnitude of the effect of higher wages on student performance mimics almost exactly the effect on teacher quality (see column (3) in Panel B of Table I).[19]

We explore further the magnitude of the effect of higher wages on student performance by using an alternative metric. The Ministry of Education classifies students into four categories according to their responses in standardized achievement tests. The lowest category correspond to students who have below basic knowledge (*Previo al inicio*), and the top category corresponding to outstanding students (*Satisfactorio*). Figure VIII displays the estimated coefficients and confidence intervals corresponding to our main specification in the sample of schools with an open contract teacher vacancy (see also Table B.11 in the Appendix). In these regressions, we use as dependent variables indicators for whether a specific student falls into one of these four categories. The results show that the effects are driven by relative changes in the two tails of the achievement distribution. The proportion of students who are below basic decreases by about 25% both in Math and Spanish in schools that receive the wage bonuses, showing that there is a strong focus on the students at the bottom of the distribution. For the case of Math test scores, there is also a large increase in the relative proportions of competent and outstanding students.

[18]As permanent positions that remain unfilled in the assignment process are later posted as vacancies for a contract teacher, the sample of schools in column (3) also includes some of the schools considered in column (4). To further isolate the effect of wage bonuses for the sample of permanent teachers, in column (3) we exclude schools that, besides having had a vacancy for a permanent teacher, also had a position only opened for contract teachers. In Appendix Table B.2 we show that there is no effect of the wage bonus on the probability that a school has an open position for permanent or contract teachers.

[19]In Appendix Figure B.6 we show that the result in column (4) is robust to different specifications. The estimated effect is stable and significant regardless of the specification for the optimal bandwidth, the variance estimation procedure, and whether regressions are defined at the school or student level.

Figure VIII: Wage Bonus and Students' Achievement Level

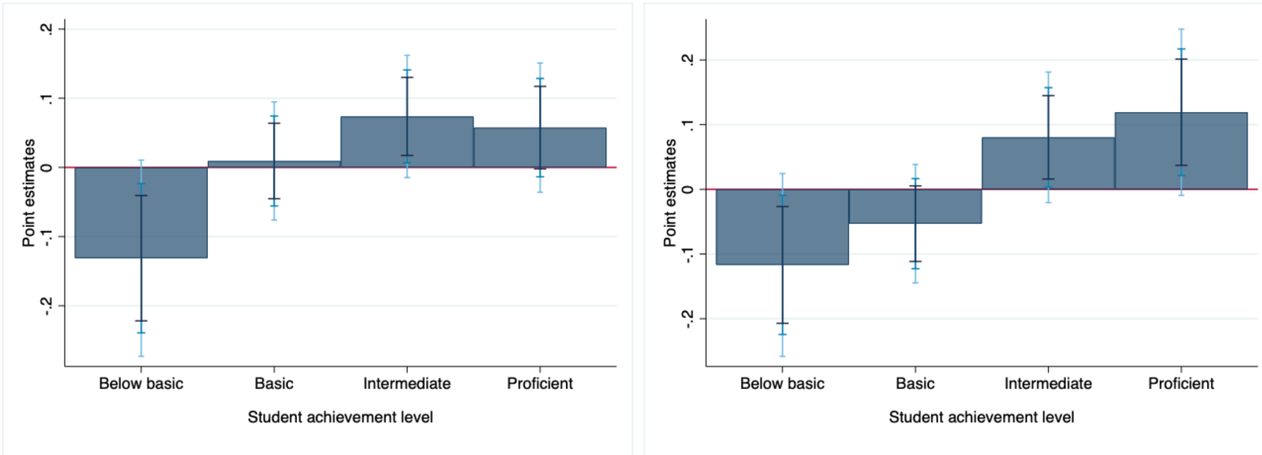

a. Spanish   *b.* Math

Notes. This table reports the effect of crossing the population threshold on student achievement in Spanish (on the left side) and Math (on the right side) classified according to four categories. These are below basic (*Previo al inicio*), basic (*En inicio*), intermediate (*En proceso*), and proficient (*Satisfactorio*). Bars and vertical lines indicates the estimated regression-discontinuity coefficients and confidence intervals (at the 90, 95 and 99% level) from a set of regression where the outcome variable is a dummy equal to one if a (fourth-grade) student falls into the corresponding category. The sample includes schools with an open position for contract teachers. Point estimates and confidence intervals are obtained using the robust estimator proposed in Calonico et al. (2014) and are reported in table format in Appendix Table B.11 In all regression, SE are clustered at the school×year level.

5 Quantifying Teacher Demand for Schools

The evidence reported in the previous section documents causal evidence that higher wages increase the number of teachers interested in working at those positions, which in turn leads to an increase in the average quality of newly recruited teachers and subsequent improvements in students' academic achievement. This shows that it is possible to undo the staggering inequality of access to quality teachers by paying more for teaching jobs in places that are less desirable.

In this section, we quantify the way teachers trade off school and local communities amenities with the compensation offered to better understand teacher preferences. We estimate an empirical model of teacher school choice using data on teachers' revealed preferences from the centralized application and assignment system. This empirical model allows us to better understand the labor supply of teachers and the reasons why the current situation is so unequal. We then use the estimated preference parameters to evaluate the implications for teacher sorting and the associated distribution of teacher quality across schools of alternative wage incentive schemes and other policies that improve other school inputs.

To better inform the empirical model of teacher school choice, we worked with the Ministry of Education in Peru to conduct a survey of teachers who participated in the job application process. Among several questions about the applications participation, the survey asked applicants to rank the most important characteristics for schools they choose to teach at (see

Table A.2 in the Appendix). Forty-four percent of teachers rank "being close to home" as the most important characteristic. The two most often cited characteristics are prestige and cultural reasons. Other characteristics that also drive decision-making for some teachers include: quality of school infrastructure, quality of students, and safety. In addition, there is clear heterogeneity in the characteristics that are most important to teachers. Moreover, teachers who score in the top quartile of the teacher evaluation test rank characteristics differently. More specifically, they appear to have stronger preferences for schools closer to home. These survey results inform the empirical model that we specify below that considers preference heterogeneity and special attention to spatial aspect of preferences including distance to *home town* and the location of the previous teaching job.

5.1 An Empirical Model of Teachers School Choice

We start by defining the utility of teacher i for being matched with school j as a function of distance, school characteristics, local amenities and compensation:

$$u_{ij} = \beta_i^c x_j^c + \beta_i^s x_j^s + \alpha_i w_j + \lambda_i d_{ij} + \epsilon_{ij} \tag{3}$$

w_j is the wage posted at school j in thousands of Peruvian Soles while x_j^c is a vector of local amenities and x_j^s a vector of school characteristics that are meant to generate variation in the individual valuations across the teaching positions. The local community characteristics include the natural logarithm of the population of the locality of the school, the time to travel (in hours) between the locality of the school and the province's capital, and a poverty index at the locality level. x_j^s includes the size of the school as measured by the natural logarithm of the number of students in the school, an indicator variable of whether the school is bilingual (Spanish and Quechua) or not, and a school input index which is the principal component of a subset of the indicator variables for whether or not the schools has access to school infrastructures (water, electricity, internet, library, room for teachers, lecture room, kitchen). Finally, we assume that ϵ_{ij} is an unobserved Gumbel idiosyncratic taste shock that is distributed *iid* across i and j with normalized scale and location.

Distance is captured by d_{ij} which is a vector of indicator variables that measure the geographic proximity between school j and the municipality of origin of teacher i as well as between school j and the previous school in which teacher i worked. We interpret both distance measures as a proxy for movement costs, which we think include both the costs of travel as well as a broader set of concerns including a preference for remaining in the school where contract teachers are located at the moment of applying for a new job.

In our model, preferences captured by $\beta_i, \alpha_i, \lambda_i$ are allowed to vary by several teacher characteristics including gender, age, experience. In addition we allow teachers who are new to the profession to have different ways of considering distance and location for one year appointments. Presumably teachers at different stages of the life-cycle may also have different costs associated with moving further away from family and social networks. We attempt to capture these considerations in a flexible model with heterogeneous teacher preferences across these dimensions.

5.2 Estimation

We observe data on teachers revealed preferences from two sources. One comes from ranked ordered lists of applications for permanent positions. The second source of data on teacher preferences comes from choices made by teachers for one year contract vacancies. The design of the assignment for permanent positions gives rise incentives for teachers to not report their preferences truthfully. Our survey of teachers asks unrestricted preferences independently of their application and the results verify that almost one third of surveyed teachers did not apply to their most preferred job posting and are providing applications that include strategic considerations. To learn about preferences from the data on ranked ordered lists, we would need to specify a model of preferences and applications, as well as school preferences. In addition, information on beliefs regarding assignment probabilities would be important as well as in Kapor et al. (2020). This would be feasible given our data on applications and survey data on preferences and beliefs. However, given the multiple sources of strategic considerations involved in the two-sided assignment mechanism used in Peru, this approach would require additional assumptions.

In this application we choose to focus on the second source of information on teacher preferences because it allows for a straightforward estimation strategy. Recall from above that within administrative units (UGEL), teachers are ranked based on their score and they are sequentially assigned to their preferred school among the ones that still have open vacancies. This procedure is iterated until all vacancies are filled and/or all teachers are assigned. In other words, conditional on the administrative unit, the assignment follows a serial dictatorship algorithm where teachers are ranked by their evaluation scores. Each teacher i is presented with a different and observable choice set Ω_i of feasible vacancies. We observe an observed choice of vacancy $\nu(i)$.

We assume that applicants know how competitive each administrative unit will be and believe that their actions have no affect on the choices of other teachers. Specifically we assume that given their known evaluation score, teachers have full information regarding the set of options they will have across all administrative areas Ω_i. This implies that, in equilibrium, teachers choose the UGEL in which their preferred feasible school is located. Given these assumptions the matching equilibrium is globally stable and preferences can be recovered from the data as the solution of a standard discrete choice model with individual-specific feasible choice sets (Fack et al. 2019).[20]

Thus our estimation strategy is to choose parameters $\{\beta, \alpha, \lambda\}$ to maximize the likelihood of observing the individual choice behaviors given by the following log-likelihood function:

$$L(\beta, \alpha, \lambda) = \frac{1}{n} \sum_{i=1}^{n} \log \frac{\exp u_{i\nu(i)}}{\sum_{j \in \Omega_i} \exp u_{ij}}, \tag{4}$$

In this model, preferences are identified if (i) teachers' test scores are independent from the taste shifters, ϵ_{ij} in (3), and (ii) the feasible choice sets are also independent from these taste shifters. The first condition is typically violated if teachers intentionally under-perform at the centralized competency test, which is unlikely to happen in this context. The second condition may not hold if there is a possibility that the decision by teacher i to accept or reject a given job posting may trigger a chain of acceptance or rejections by other teachers which may feed back into teacher i's set of feasible school alternatives (Menzel 2015). Preference cycles of this sort are ruled out in our setting since school preferences are homogeneous, which imply monotonicity along the applicants' (score-based) ranking with respect to any possible chain of acceptances or rejections.

In estimation we use the same teachers that participate in the assignment process of short-term contract teaching jobs. We use the same sample both for estimation and counterfactual analysis. To illustrate the role of heterogeneous preferences, we augment the estimates of the baseline model for various groups/types of teachers. In the counterfactual analysis we use the estimates of the most flexible model with interaction terms and polynomials for the different teachers' types in order to better fit the choices observed in the data.

5.3 Model Estimation Results

Table III reports selected model estimates for wages, those for the proxies of local community amenities, an index of school infrastructure and school inputs, as well as interaction effects with teachers' ethnolinguistic traits and the racial composition of the district where the school is located. Column (1) reports the average preference parameters while Columns (2)-(5) show how these parameters vary across different types of teachers according to their age, location of origin, and gender. Teachers value compensation although men are significantly more elastic to wages. Teachers prefer schools that are larger, are in less poor areas and schools that have more infrastructure and educational inputs. Urban teachers dislike schools that are located in poorer areas. Teacher-school match effects are very strong, especially for minorities who have a clear

[20] It should be noted that the majority of the applicants in our sample select the administrative unit where they currently work and/or where they reside (84% in the recruitment drive of 2015 and 86% in 2017).

Table III: Model Estimates – Selected Preference Parameters

		Heterogeneity			
	All	Base	Age < 30	Rural	Male
Wage	0.635***	0.111	0.243	0.134	0.983***
	(0.0657)	(0.105)	(0.205)	(0.146)	(0.137)
Poverty Score Locality	-0.0817***	-0.108***	-0.0201	0.0704***	0.0537***
	(0.00750)	(0.0106)	(0.0224)	(0.0182)	(0.0162)
School Infrastructure Index	0.0177	0.0376**	0.0208	-0.00236	-0.0514*
	(0.00975)	(0.0137)	(0.0281)	(0.0233)	(0.0209)
Locality Infrastructure Index	0.0510***	0.0397	0.0520	0.0128	-0.0225
	(0.0148)	(0.0212)	(0.0440)	(0.0350)	(0.0313)
Match effects					
Mestizo × Share Mestizo in school's district	0.460***	0.296***	-0.393*	0.405**	0.232
	(0.0619)	(0.0877)	(0.176)	(0.154)	(0.131)
Quechua × Share Quechua in school's district	1.670***	2.416***	-0.125	-1.538***	-0.306
	(0.105)	(0.173)	(0.313)	(0.217)	(0.212)
× Bilingual school	1.106***	1.269***	0.0686	-0.159	-0.291*
	(0.0654)	(0.103)	(0.200)	(0.144)	(0.136)
Aimara × Share Aimara in school's district	2.338***	3.560***	-2.547***	-0.159	-0.649
	(0.280)	(0.449)	(0.632)	(0.611)	(0.550)
× Bilingual school	0.894***	0.959***	0.799	-0.214	-0.202
	(0.167)	(0.259)	(0.460)	(0.372)	(0.340)

NOTES. Standard errors in parenthesis. *** $p < 0.001$, ** $p < 0.01$, and * $p < 0.05$. This table displays the estimates of the model for teachers' preferences described in Section 5. We use the 8,190 teachers assigned in 2015 along with the 10,569 teachers assigned in 2017 and consider the two samples as independent cross sections. We construct the feasible choice sets by first determining the score of the lowest ranked applicant in each school (which we will call cutoffs). If the school hasn't filled all vacancies it is feasible by definition, if the school is full, it is feasible only if the teacher has a score above the cutoff. We then estimate this discrete choice model with personalized choice sets by maximum likelihood. Across both years, the feasible choice sets contain 5,672 schools on average.

preference for schools in specific areas that share their ethnolinguistic traits and particularly so if schools are bilingual.

The full set of estimated preference parameters across several model specifications are presented in Table C.1 in the Appendix and include a granular set of distance dummies as well as the full set of locality characteristics that determine the wage schedule for public-sector teachers in Peru. One notable result is that distance of a job opening to the teachers home town or their current job are both very important. To visualize the importance of distance in explaining teachers' choices over schools, Figure IX plots the implied wages needed to compensate teachers from moving far away from where they live or from where they previously worked. For instance, it would take a wage that is approximately 7 times higher than the current wage in order to make teachers indifferent between working in the school where they currently are located and another school situated 100 kilometers away. To the extent that higher-quality teachers are mostly located in urban areas, as shown in Figure II, public policies aimed at enhancing the local supply of teachers in remote areas might be a promising alternative to wage incentives in order to reduce regional inequalities in the quality of teachers.

Figure IX: Estimated Preference for Distance

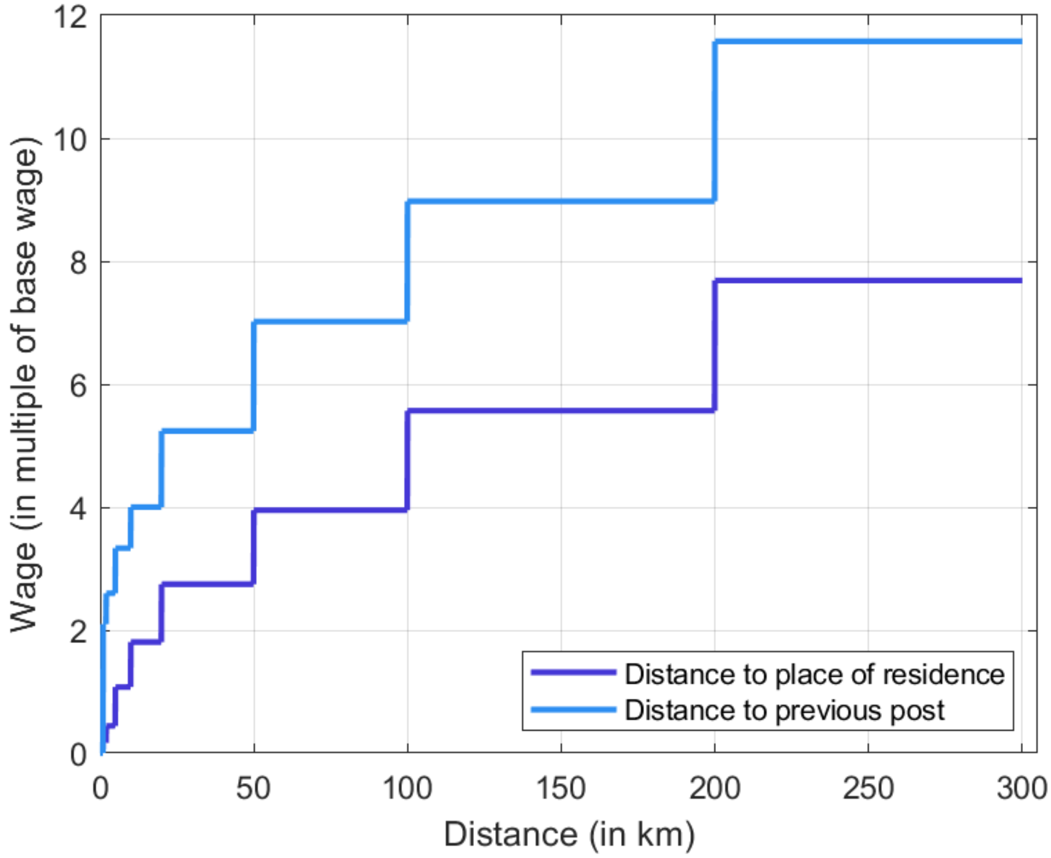

Notes. This figure draws the indifference curves of teachers on the wage-distance axis using the two definitions of distance (from the municipality of origin or from the previous job). Distance is measured in km and wages are measured in multiples of the base wage (which is 1555 soles).

5.4 Visualizing Fitted Utility Across Space

The spatial aspect to teachers preferences is a very salient feature of the current wage schedule, making jobs at schools that are farther away and have smaller populations, have higher wages. To describe the effect of the current wage policy, we visualize the effect of wage bonuses on the population and distance space that determines the wage bonuses. Specifically we take the estimated parameters and evaluate the utility for each teacher at each school with and without the wage bonuses. The left panel of Figure X shows an estimate of the average median utility at each combination of population and distance to provincial capital. The darker colors in the top left corner of the figure show lower levels of utility are associated with vacancies located in localities with very low population and also are multiple hours away from the provincial capital. The lighter colors in the bottom center of the figure show that schools located closer to the provincial capital and are in localities with 500-1000 people tend to provide more utility

Figure X: Change in Fitted Teacher Utility With the Current Wage Bonus

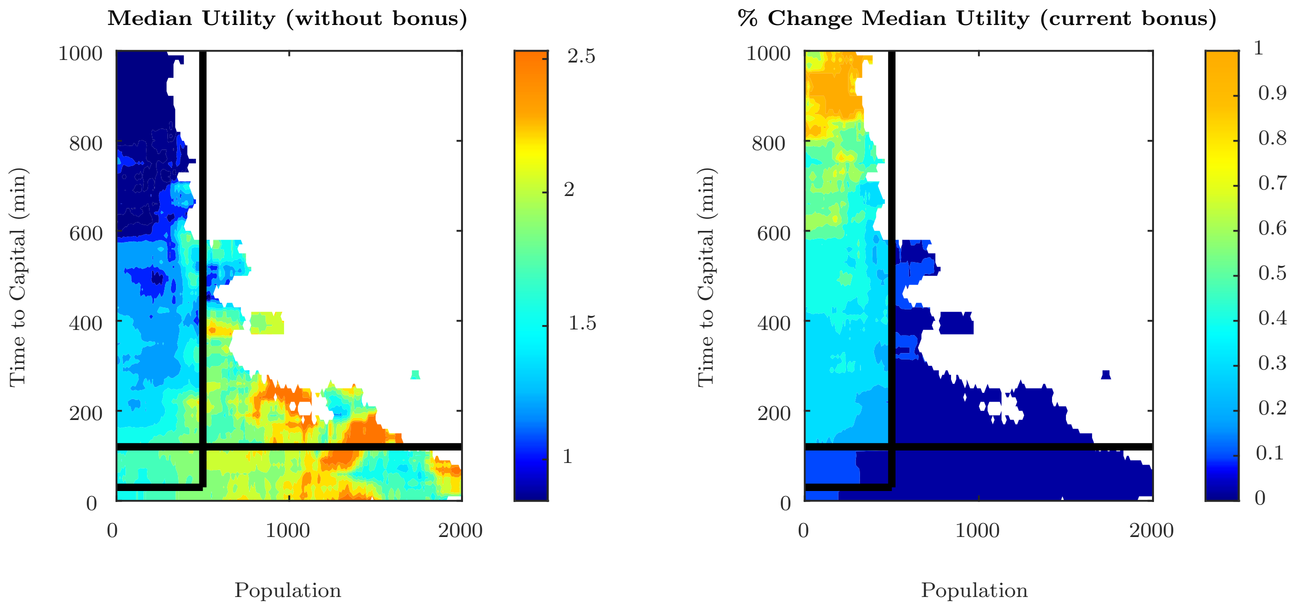

NOTES: The left panel shows the average median utility associated with vacancies at each combination of distance to the capital and the population. The right panel shows the percent difference in average median utility when the current wage bonus policy is implemented.

to potential teachers. Given that the wage bonuses are not included, it is not surprising to see the utility is relatively continous around the eligibility cutoffs of 120 mins and 500 population as well as other thresholds.

The right panel of Figure X shows the percent change in average median utility when the current wage bonus creates. Here the wage bonus is included and the area that determines the highest wage bonus of 500 soles, shows a striking change in the percent change in average median utility. Areas with lower wage bonuses seem to have increments of no more than 10% while the extremely rural vacancies see increments of over 30%. This is consistent with the data we observe from the application data that shows that at the threshold, applications show the wage bonus of 500 had a large effect on how desirable the options were.

5.5 Model Fit

In order to bolster the credibility of the model-based counterfactual experiments discussed in the next section, we first assess how well the model predicts some key moments in the data. In particular, it is important to corroborate the validity of the estimated wage elasticity estimated off the observed cross-school variation in posted wages. To do so, we check the consistency between the model estimates and the RD estimates presented above. Provided that only wages change at the population cutoff, the estimated size of the jump in teacher scores may be used for model validation. We thus simulate teachers' choices using the estimated preference

Figure XI: Comparing Simulated Threshold Crossing Effects at the Population Cutoff)

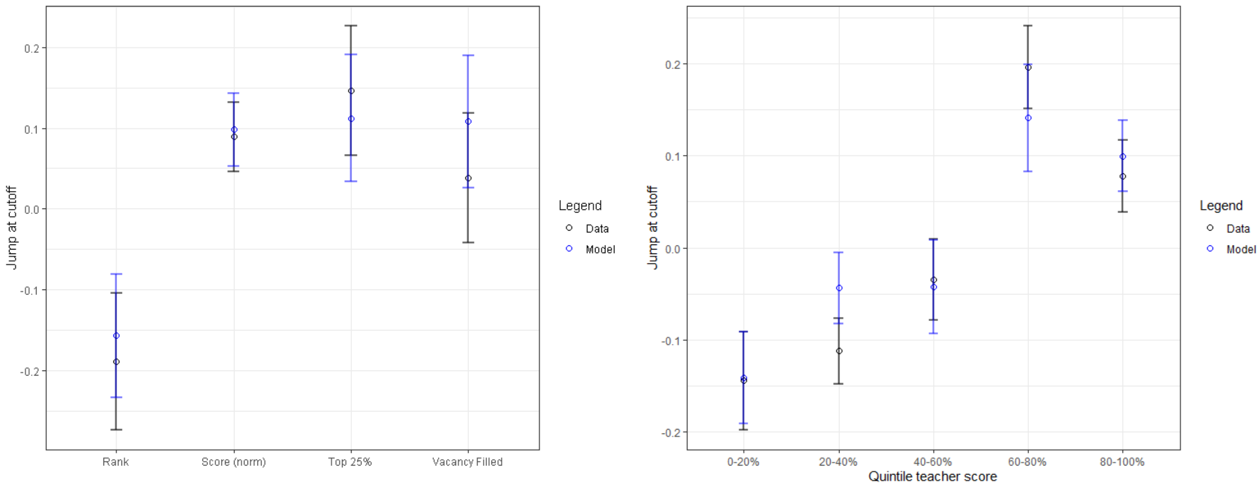

a. Model Validation: Main Effects

b. Model Validation: By Teachers' Quality

Notes. To assess model fit, we predict indirect utilities for each teacher and simulate the match using the serial dictatorship algorithm. We then recompute the jump in teachers' test scores at the population cutoff and compare it with the estimated jump observed in the data.

parameters, replicate the RD analysis on simulated data, and compare the resulting estimates across models and data. Figure XI shows the result of this exercise. The estimated model seems to predict very well the different selection patterns triggered by the wage bonus that we observe in the data, both for the main effects (Panel A) as well as the compositional changes along the distribution of the teachers' competency score.

5.6 Adjusting Compensation to Provide Equal Opportunities

We use the estimated preference parameters to predict the counterfactual wage bonus that would be sufficient to attract one teacher who is above the median of the score distribution in each teaching vacancy made available through the 2015 and 2017 centralized recruitment drives. To do so, we compute for each vacancy that was filled with a teacher below the median of the score distribution the implied wage distribution that would make any above-median teachers indifferent between their current match and the match with the school associated to that vacancy. We then take the minimum of this distribution for each school, which allows us

Figure XII: Compensation to Equalize Expected Utility of Prospective Teachers

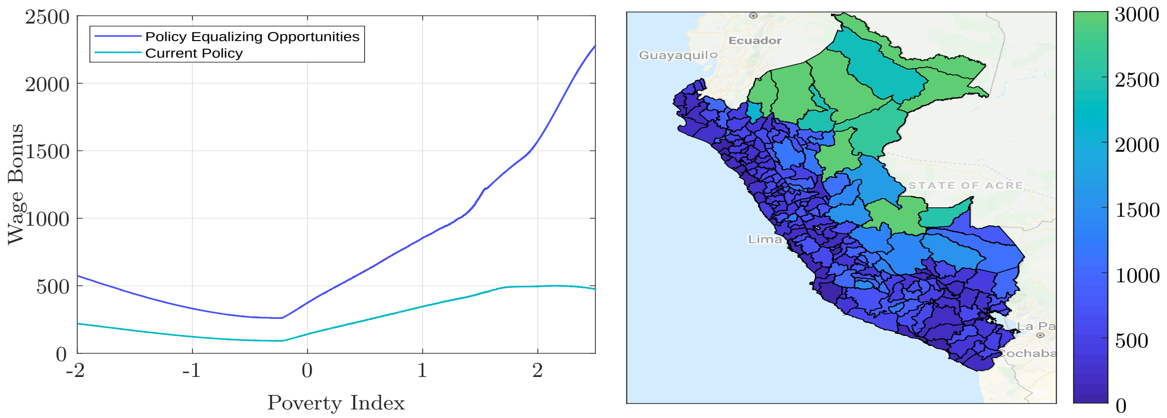

a) Wage Schedule That Equalizes Access **b) Equalizing Access Wages by Region**

NOTES: The figure in the left panel shows the wage bonus from current policy and wage bonus we calculate is needed based on the parameter estimates. The right panel shows the wage bonus needed on average by geographic location.

to compute the resulting average wage bonus across vacancies at the province level.[21]

The right panel of Figure XII shows the geographic distribution of the wage incentives that would eliminate the observed regional inequality in teachers' quality. While it is clear that the current policy is progressively targeting more disadvantaged locations, the magnitude of the monetary incentives in place seems to fall short in driving the extent of teachers sorting by quality across space. The left panel of Figure XII further illustrates this point by showing a high correlation between the counterfactual wage bonuses and poverty across provinces (upper panels). While there is a mildly positive correlation between wages and poverty in the current policy, it becomes much steeper after the mean level of the poverty index with the counterfactual policy.

Figure XIII shows the CDF of the implied minimum wage bonus across school and locality characteristics, where the status quo bonus is indicated with a vertical (red) bar. For instance, under the current policy 80% of the teaching vacancies are filled with relatively high-quality teachers in schools that are located in provincial capitals, whereas it would take a wage bonus that is eight times larger than the status quo bonus in order to accomplish the same objective in localities that are 5 hours away or more from provincial capitals (middle panel). Similarly, 60% of the vacancies are filled with relatively high-quality teachers under the current policy in schools that are above-median in the distribution of students' test score in mathematics,

[21]This exercise has two drawbacks. First, we don't take into account that the policy under study may alter the equilibrium cutoffs, which would in turn affect the choice sets of teachers. Indeed, we might attract an above median teacher from a school that would in the end be left with a lower-quality teacher. Second, by taking the minimum of the wage bonus distribution for each vacancy we might end up considering the same teacher more than once. If this was the case, the resulting total wage bill would be a lower bound of the actual cost of the policy. Notice though that even when using this restrictive criterion, 3,792 of the 8,581 selected teachers are different individuals.

Figure XIII: Cumulative Distribution of Counterfactual Wage Bonuses Needed

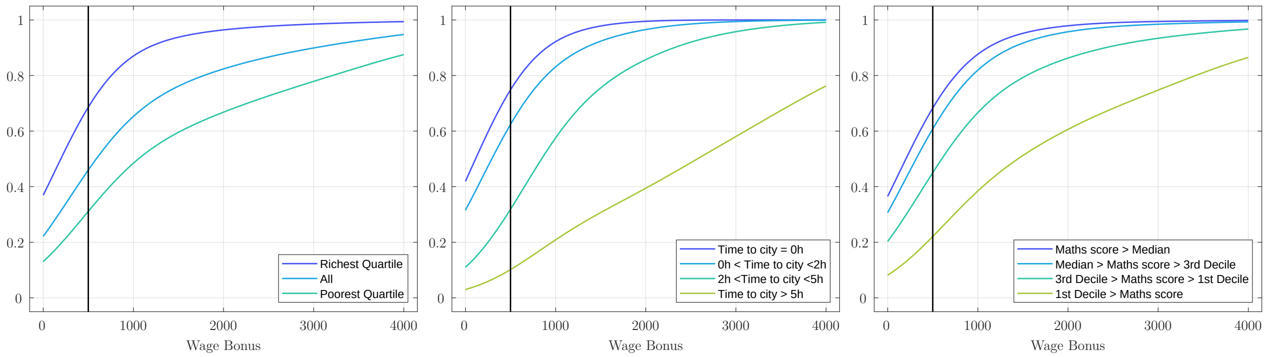

Poverty Status **Remoteness** **Students' Math Achievement**

NOTES. The left panel plots the cdf of the counterfactual wage bonuses depending on the poverty index of the municipality in which the school is located. The right panel looks at the heterogeneity with respect to the time (in hours) it takes to travel from the school to the largest city of the province. The bottom panel looks at the heterogeneity with respect to the score students achieved at the ECE maths test.

whereas it would take a wage bonus 6 times larger to do so in schools at the bottom decile of math scores (right panel).

5.7 Quantifying the Inequality Due to School Inputs and Community Amenities

School inputs and local community amenities contribute beyond the remoteness of the location to making some schools less desirable to potential teachers. To gauge how relevant these other characteristics we compare the wage bill from our prior benchmark, we compare wage schedules after eliminating other structural inequalities. Table IV provides a breakdown of both the total monthly wage bill implied by the policy equalizing teaching quality, which is three times higher than the current policy, under alternative policy levers that may partly contribute to achieve the same objective in terms of the distribution of teaching quality. We first investigate what would be the effect of removing all the observed structural inequalities between localities. This would only allow to save about 20% on the total wage bill of the policy that would attract an above median teacher in every school. Given how much distance matters in teachers choices (see Figure IX) a promising alternative may by investing in local teaching quality. We thus simulate such a policy by artificially setting all distances to zero between an above median teacher and all schools from the same province. This would reduce the total wage bill of the policy by 30%.

Table IV: Alternative Policy Counterfactuals

	Total Wage Bill per Month		Net Present Value	% of Total Cost
	2015	2017		
Policy 1: Wage Bonus Only				
Cost of Current policy (benchmark)	3.53	4.66	2,264.97	35.93
Cost of Equalizing teacher quality	12.33	13.30	6,303.71	100
Policy 2: Equalizing Structural Inequalities				
Infrastructure	0.26	0.21	127.88	2.03
Time to travel	1.07	0.01	356.43	5.65
Size school	0.36	0.82	481.92	7.65
Poverty index	0.01	-0.01	23.65	0.37
Village Population	0.01	0.12	-121.59	-1.93
Bilingual Schools	0.78	0.25	295.61	4.69
Adjusted Cost of Equalizing teacher quality	9.75	11.90	5,139.81	81.54
Policy 3: Increasing Local Supply of High-quality Teachers				
Adjusted Cost of Equalizing teacher quality	7.16	10.72	4,377.67	69.45

NOTES. The table displays the total cost per month (in millions of soles) of attracting an above median teacher in each vacancy for each year (column 1 and 2). It also shows the net present value of each policy on a duration of 20 years using a discount factor of 2% (column 3). The top panel shows this cost when wage incentives are the only policy instrument available. In the second panel of the table we remove inequalities in schools/locality characteristics one by one and recompute the wage bonuses needed to equalize opportunities. In the bottom panel we simulate local increases of teaching quality by setting the distance to zero between an above median teacher and all schools within the same province. We then recompute and display the wage bonuses needed to equalize opportunities.

6 Equilibrium Counterfactuals

While the analysis done in the previous section may help the reader to get a sense of what the level of compensation would need to be to attract good quality teachers to remote schools, it does not take into account that, in equilibrium, teachers would sort in a way that might be very different than what can be predicted ex-ante. Indeed, in the previous exercise, the same teacher could be selected more than once to fill two different schools and thus does not take into account that high quality teachers are scarce. The wage schedule derived is thus a lower bound on what it would need to be to actually reshuffle teacher quality in equilibrium. We thus propose in this section a procedure that would give us the cost efficient wage schedule that would increase demand for schools in underprivileged areas such that, in equilibrium, every school in the country would get assigned at least one above median quality teacher.

6.1 Methodology

Let us consider a framework where we have a set of teachers T, a set of schools S and a set of possible wages W. We define a contract as a school, a teacher and a wage where the set of possible contracts is denoted $X = T \times S \times W$. We know from Hatfield & Milgrom (2005) that, under some regularity conditions on preferences, a stable set of contracts always exist. This means that there exists an allocation such that schools and teachers would not want to deviate and break their current match for any other agent at the proposed wages. In a traditional labor market, this framework would typically be used in order to take into account that other dimensions such as wages or amenities can be leveraged as an additional market clearing device. However, in most labor markets for public servants, wages usually act as a tool for the central planner to achieve a given social objective. There is thus scope for using this framework to guide policy makers in finding the most cost effective way to achieve such an objective. We will thus show that, by encoding a given social objective in schools' preferences, the matching with contracts framework can indeed be used to find the most cost effective stable set of contracts which would reach this objective.

We first start by defining agents' preferences over contracts. Teachers' preferences over each school-wage pair were already identified and estimated in the previous section. We thus need to specify schools' preferences such that the school proposing matching with contracts algorithm will give us the most cost effective wage schedule that would reach our social objective under the assignment mechanism currently in place in Peru. We define our social objective as reaching an allocation under the current assignment mechanism in place in Peru where every school would have at least one teacher with a score above a given threshold $\bar{s}$. We thus assume that all schools rank contracts for their first open slot such that:

- $\{(i, w_1) : s_i \geq \bar{s}\} \succ \{(l, w_2) : s_l < \bar{s}\} \quad \forall (i, l, w_1, w_2) \in T^2 \times W^2$
- $\{(i, w_1) : s_i \geq \bar{s}\} \succ \{(l, w_2) : s_l \geq \bar{s}\} \quad \forall (i, l) \in T^2$ and for any $w_1 < w_2$
- $\{(i, w_1) : s_i < \bar{s}\} \succ \{(l, w_2) : s_l < \bar{s}\} \quad \forall (i, l) \in T^2$ and for any $w_1 < w_2$
- $\{(i, w)\} \succ \{(l, w)\} \quad \forall (i, l, w) \in T^2 \times W \iff s_i > s_l$

The first requirement states that a school would prefer any teacher with a score above $\bar{s}$ to a teachers below $\bar{s}$ irrespective of the wage. This makes sure that schools will increase wages until at least one good quality teacher is willing to accept their offer. The second and third requirement state that among teachers with a score above $\bar{s}$ and among the teachers below $\bar{s}$ schools would always prefer to hire at the cheapest cost. The fourth requirement allows to break ties by stating that for a given wage, schools would prefer the highest quality teacher. This last requirement also makes sure that the final allocation can be reached by using the same assignment mechanism as the one currently used in Peru. For the remaining slots, we assume that schools instead use the following ranking:

- $\{(i, w_1)\} \succ \{(l, w_2)\} \quad \forall (i, l) \in T^2$ and for any $w_1 < w_2$
- $\{(i, w)\} \succ \{(l, w)\} \quad \forall (i, l, w) \in T^2 \times W \iff s_i > s_l$

This makes sure that schools will stop increasing wages to compete for good quality teachers once they managed to attract one. Given that good quality teachers are scarce, schools would never stop increasing wages without this requirement and the algorithm will never converge. Of course this can be adjusted depending our the objective function. If we want to attract at least two teachers with a score above $\bar{s}$ in every school, then we need to assume that schools have the preferences described in the first place for their first two slots. We can now state our main result.

Proposition 6.1

(i). Under the preferences described above, the outcome of the school-proposing matching with contracts algorithm gives the most efficient wage schedule that achieves the social objective of attracting at least one teacher with a score above $\bar{s}$ in every school.
(ii). Each iteration of the algorithm gives the allocation that minimizes the share of unfilled schools with at least one teacher with a score above $\bar{s}$ under the budget constraint of the sum of the wages proposed at this round.

The outcome of this algorithm is thus giving us the minimum wage that would attract a good quality teacher in each school while taking into account equilibrium sorting. On top of that, each step of the algorithm gives us the allocation that minimizes the share of unfilled schools under the budget constraint of the sum of the wages proposed at this round. To see that, consider running the algorithm by imposing the additional stopping rule "stop when the sum of the wages proposed exceeds x", then the outcome of the algorithm would give us the optimal wage schedule that minimizes the share of unfilled school given this budget constraint.

6.2 Algorithm

To give a sketch of proof for Proposition 6.1, let us first describe the school-proposing matching with contracts algorithm:

Round 1: Each school proposes to its most preferred teacher-wage pair. Teachers are tentatively assigned to the proposing schools at the wage specified in the contract. All schools which did not fill at least one vacancy move to the next round.

Round k: Each unassigned school proposes to its next preferred teacher-wage pair. Teachers choose their preferred offer from those made in all rounds up to k. All unfilled schools move to the next round.

The algorithm stops once all schools are filled or once all unfilled schools run out of offers. However, given our assumptions on schools' preferences, we can show that we can rewrite this algorithm and simplify it significantly. Indeed, given that all schools have the same preferences and that, for a given wage, teachers with a score below $\bar{s}$ are dominated by teachers with a score above $\bar{s}$, we can decompose the algorithm in two stages. First, start by running the algorithm only with above $\bar{s}$ teachers. Then run the serial dictatorship algorithm with the remaining slots and the remaining teachers using the wages resulting from the first round. We thus describe here the algorithm which allocates the n teachers which have a score above $\bar{s}$.

Round 1: Each school proposes to the highest quality teacher at the lowest wage possible. This teacher is tentatively assigned to its preferred school. All schools which still have an unfilled vacancy move to the next round.

Round n: Each school with remaining vacancies proposes to the lowest quality teacher at the lowest wage possible. This teacher is tentatively assigned to its preferred school. All schools which still have an unfilled vacancy move to the next round.

Round $n+1$: Each school with all vacancies empty start proposing to the highest quality teacher at a slightly higher wage. This teacher chooses its preferred offer from those made in all rounds up to $n+1$. All schools which still have an unfilled vacancy move to the next round.

Round $k > n+1$: Each school with all vacancies empty start proposing either to their next preferred teacher at the same wage or to the highest quality teacher at a slightly higher wage. Teachers choose their preferred offer from those made in all rounds up to k. All schools which still have an unfilled vacancy move to the next round.

The algorithm stops once every school has at least one of its slot filled. We can also show that this algorithm is equivalent to iterating the serial dictatorship algorithm taking teachers' score as priorities and only letting the unfilled schools increase their proposed wages at each round which makes it very easy to implement (put proof in Appendix?).

6.3 Results

From the description of the algorithm, we can easily see that at each iteration, the sum of the wages proposed is weakly increasing while the share of unfilled school is weakly decreasing. We can thus use that on top of Proposition 6.1 to draw a cost efficiency frontier showing us the cost of the most efficient policy that would lower the share of unfilled school to a given level.

Figure XIV: The Cost-Efficiency Frontier and the Current Wage Bonus Policy

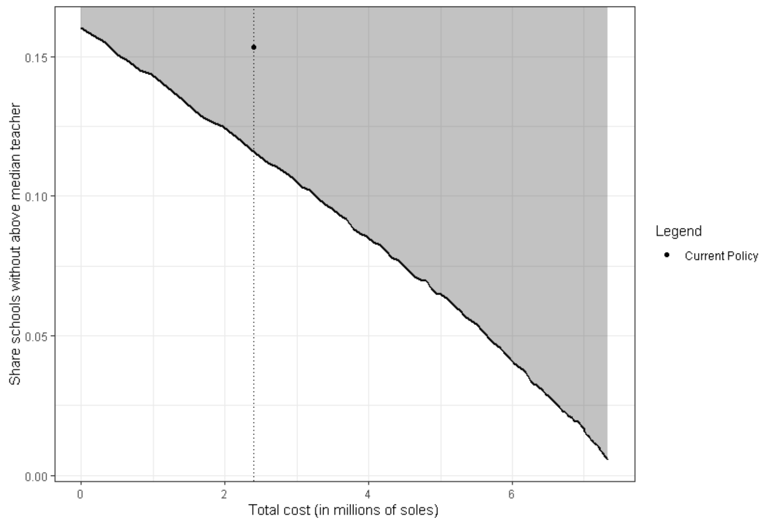

Figure XIV draws this cost efficiency frontier and shows that the policy in place is far from being optimal. This suggests that a better targeted and informed policy would perform significantly better at reducing spatial inequalities in access to teaching quality.

Figure XV depicts what this optimal schedule would be and shows that the current wage schedule overcompensates some areas while under-compensating others. This suggests that the current schedule is not taking into account that high quality teachers are more scarce in some areas than others and that this scarcity is not fully explained by the categories currently used to determine the wage bonuses. To confirm this intuition, Figure XVI shows that it would be actually more efficient to compensate less the current rurality categories and give bonuses to single teachers schools or bilingual schools which are more likely to be located in extremely remote areas where good quality teachers are scarce. We next show that by changing the population and time cutoffs that define the rurality categories, we are able to get closer to the optimal wage schedule. Looking at the R^2 of a regression of the optimal wage schedule on this categories allows us to measure to which extent the categories we propose are able to fit this schedule. We can see that by increasing the population and time threshold for the Rural 1 categories, we are able to significantly increase the R^2 of this regression suggesting that a redefinition of these categories would allow to get closer to the cost efficiency frontier.

Figure XV: The Optimal Wage Schedule 1

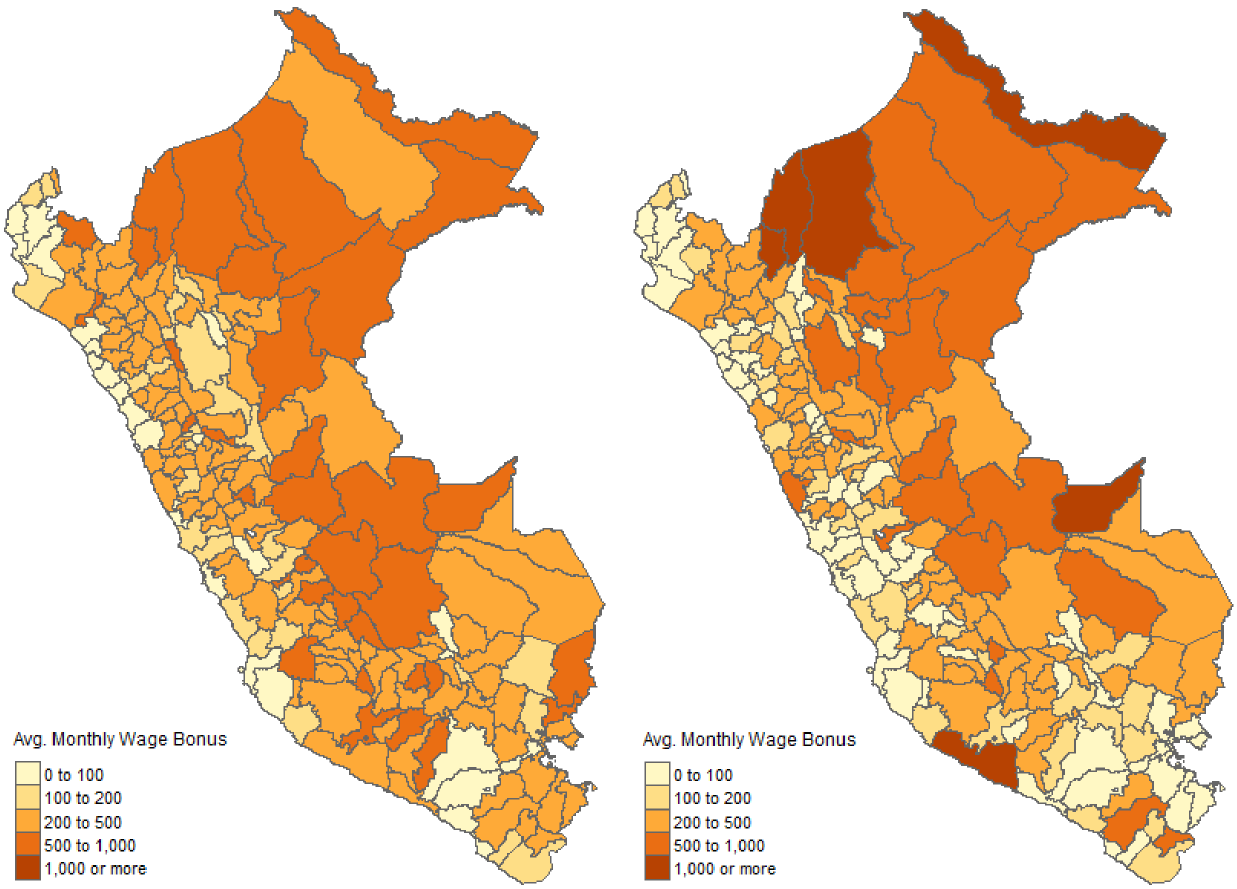

Current Policy **Optimal Policy**

Figure XVI: The Optimal Wage Schedule 2

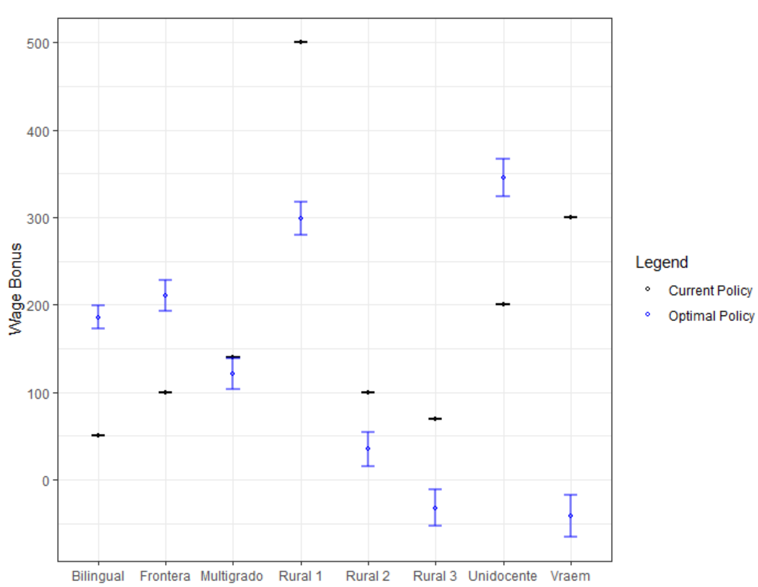

7 Conclusion

This paper studies how increasing teacher compensation at hard-to-staff schools can reduce structural inequality in the access to high-quality teachers. Using rich administrative data from Perú, we document dramatic inequities in schooling inputs and teacher quality to which students have access. This is particularly worrying given the evidence that teacher quality has long term consequences on adult labor market outcomes and can thus perpetuate the initial inequality students face.

We show how teacher compensation can contribute to reducing the inequality in educational opportunities offered to students from poorer rural locations in Peru. The Peruvian education context is uniquely well suited to study this question for three reasons. First, the government implementated a policy that generated arbitrary cutoff rules for wage increases that allow for a credible empirical strategy built around a crisp regression discontinuity design. Second, the entire public school system organizes teacher job postings, teacher job applications and final assignments in a centralized way, providing rich data on the entire process through which a teacher is assigned to a particular post. This system also provides an internally consistent measure of teacher quality that is specific to the job. Third, the large presence of contract teachers that are assigned to temporary teaching positions creates built-in flexibility in the teacher labor market, which in turn can generate large sorting responses to wage incentives within a relatively short time span.

We use the data and policy variation to conduct two sets of empirical excersizes. The first is to show causal evidence that increasing teacher pay has both recruitment and productivity effects. Specifically we find that unconditional wage increases are successful in effectively attract and retain talent to public schools. These higher wages also cause significantly higher retention rates when combined with transparent, merit-based assignment rules for contract teachers. We are further able to look at the productivity effects of these newly recruited workers, and document that students in high wage schools perform better in standardized tests. The observed increase in productivity is highly correlated with the increase in average teacher talent across schools. In fact, the policy effect on student outcomes is entirely driven by students in schools that had multiple openings during the period when the policy was in place, while it is estimated to be a tight zero in schools where no new openings were available.

The second empirical contribution is to quantify the way teachers trade off wages and local school and community amenities leveraging the rich data on applications and job postings from the centralized assignment system. The model estimates nicely replicate the reduced-form findings and they reveal a strong dis-utility effect of distance from teachers' residence to the school, which is much larger in magnitude than the estimated wage elasticity. Counterfactual changes in the wage bonus aimed at reducing the extent of cross-regional inequality in the quality of teachers are predicted to be very large, and they should be probably accompanied by complementary policy interventions in order to accomplish the stated objectives in a more cost-effective way.

References

Acemoglu, D., Johnson, S., and Robinson, J. A. (2001). The colonial origins of comparative development: An empirical investigation. *American Economic Review*, 91(5):1369–1401.

Araujo, M. C., Carneiro, P., Cruz-Aguayo, Y., and Schady, N. (2016). Teacher quality and learning outcomes in kindergarten. *The Quarterly Journal of Economics*, 131(3):1415–1453.

Artiles, M. (2020). Within-group heterogeneity in a multi-ethnic society. Technical report, Working Paper, Mimeo.

Barrera-Osorio, F. and Raju, D. (2017). Teacher performance pay: Experimental evidence from Pakistan. *Journal of Public Economics*, 148(C):75–91.

Bau, N. and Das, J. (2020). Teacher value added in a low-income country. *American Economic Journal: Economic Policy*, 12(1):62–96.

Berlinski, S. and Ramos, A. (2020). Teacher mobility and merit pay: Evidence from a voluntary public award program. *Journal of Public Economics*, 186:104186.

Bold, T., Filmer, D., Martin, G., Molina, E., Stacy, B., Rockmore, C., Svensson, J., and Wane, W. (2017). Enrollment without learning: Teacher effort, knowledge, and skill in primary schools in africa. *Journal of Economic Perspectives*, 31(4):185–204.

Buhl-Wiggers, J., Kerwin, J., Smith, J., and Thornton, R. (2017). The impact of teacher effectiveness on student learning in africa.

Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust Nonparametric Confidence Intervals for Regression?Discontinuity Designs. *Econometrica*, 82:2295–2326.

Calonico, S., Cattaneo, M. D., and Titiunik, R. (2015). Optimal Data-Driven Regression Discontinuity Plots. *Journal of the American Statistical Association*, 110(512):1753–1769.

Card, D. and Krueger, A. B. (1992a). Does school quality matter? returns to education and the characteristics of public schools in the united states. *Journal of political Economy*, 100(1):1–40.

Card, D. and Krueger, A. B. (1992b). School quality and black-white relative earnings: A direct assessment. *The Quarterly Journal of Economics*, 107(1):151–200.

Cattaneo, M. D., Jansson, M., and Ma, X. (2020). Simple local polynomial density estimators. *Journal of the American Statistical Association*, 115(531):1449–1455.

Chetty, R., Friedman, J., and Rockoff, J. (2014a). Measuring the impact of teachers ii: Teacher value-added and student outcomes in adulthood. *American Economic Review*, 104(9):2633–2679.

Chetty, R., Friedman, J. N., and Rockoff, J. E. (2014b). Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood. *Forthcoming. American Economic Review*.

Dal Bo, E., Finan, F., and Rossi, M. A. (2013). Strengthening state capabilities: The role of financial incentives in the call to public service. *The Quarterly Journal of Economics*, 128(3):1169–1218.

de Ree, J., Muralidharan, K., Pradhan, M., and Rogers, H. (2018). Double for Nothing? Experimental Evidence on an Unconditional Teacher Salary Increase in Indonesia. *The Quarterly Journal of Economics*, 133(2):993–1039.

Dell, M. (2010). The persistent effects of peru's mining mita. *Econometrica*, 78(6):1863–1903.

Deserranno, E. (2019). Financial incentives as signals: Experimental evidence from the recruitment of village promoters in uganda.

Deserranno, E., Kastrau, P., and Leon-Ciliotta, G. (2021). Promotions and productivity: The role of meritocracy and pay progression in the public sector. Technical report, Working paper.

Duflo, E., Dupas, P., and Kremer, M. (2015). School governance, teacher incentives, and pupil–teacher ratios: Experimental evidence from kenyan primary schools. *Journal of Public Economics*, 123:92–110.

Fack, G., Grenet, J., and He, Y. (2019). Beyond Truth-Telling: Preference Estimation with Centralized School Choice and College Admissions. *American Economic Review*, 109(4):1486–1529.

Finan, F., Olken, B. A., and Pande, R. (2017). The personnel economics of the state. In Banerjee, A. and Duflo, E., editors, *Handbook of Field Experiments, Volume II*. North Holland.

Fryer, R. G. (2013). Teacher Incentives and Student Achievement: Evidence from New York City Public Schools. *Journal of Labor Economics*, 31(2):373–407.

Gallegos, S., Neilson, C., and Calle, F. (2019). Screening and recruiting talent at teacher colleges using pre-college academic achievement.

Gerard, F., Rokkanen, M., and Rothe, C. (2020). Bounds on treatment effects in regression discontinuity designs with a manipulated running variable. *Quantitative Economics*, 11(3):839–870.

Gregorio, S. D., Gallegos, S., and Neilson, C. (2019). Decomposing teachers contributions to learning. Technical report, Working Paper, Mimeo.

Kapor, A. J., Neilson, C. A., and Zimmerman, S. D. (2020). Heterogeneous beliefs and school choice mechanisms. *American Economic Review*, 110(5):1274–1315.

Menzel, K. (2015). Large matching markets as two-sided demand systems. *Econometrica*, 83(3):897–941.

Mizala, A. and Ñopo, H. (2016). Measuring the relative pay of school teachers in latin america 1997–2007. *International journal of Educational development*, 47(1):20–32.

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

Rosen, S. (1986). Prizes and Incentives in Elimination Tournaments. *American Economic Review*, 76(4):701–715.

Rubin, D. B. (1986). Statistics and causal inference: Comment: Which ifs have causal answers. *Journal of the American Statistical Association*, 81(396):961–962.

A Appendix - Descriptive Statistics

Table A.1: School Inputs and Outputs

	Extremely rural schools		Rural schools		Urban schools	
	Mean	Sd	Mean	Sd	Mean	Sd
Teaching Staff						
Teacher score in 2016 test (std)	-0.21	0.87	-0.02	0.88	0.29	0.83
Teachers with permanent contract (%)	0.44	0.34	0.64	0.25	0.76	0.17
Teachers with temporary contract (%)	0.43	0.35	0.24	0.23	0.18	0.17
Infrastructure						
No water	0.36	0.48	0.15	0.36	0.20	0.42
No electricity	0.28	0.45	0.06	0.23	0.10	0.32
Sport facility	0.15	0.36	0.41	0.49	0.50	0.53
Community Characteristics						
Sewage in town/village	0.23	0.42	0.49	0.50	0.40	0.52
Doctor in town/village	0.39	0.49	0.66	0.47	0.50	0.53
Library in town/village	0.01	0.10	0.05	0.22	0.20	0.42
School Size						
Single-teacher school	0.12	0.32	0.02	0.15	0.01	0.07
Multigrade school	0.78	0.41	0.46	0.50	0.02	0.14
Number of students	57	35.63	100	76.09	397	265.58
Number of teachers	3.19	1.89	6.25	3.87	19.19	11.57
Number of schools	1773		3130		3367	

Figure A.1: Correlation Between Teacher and Student Scores

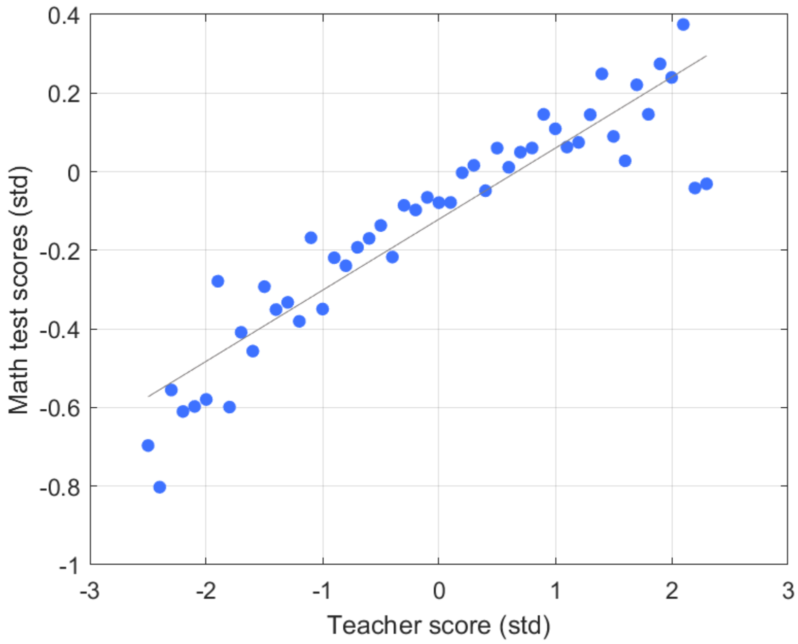

NOTES. This figure depicts the relationship between teacher competency scores and student test scores. Both measures are the residuals from linear regressions of standardized scores on school and grade fixed-effects.

Figure A.2: Contract teacher selection (Low wage)

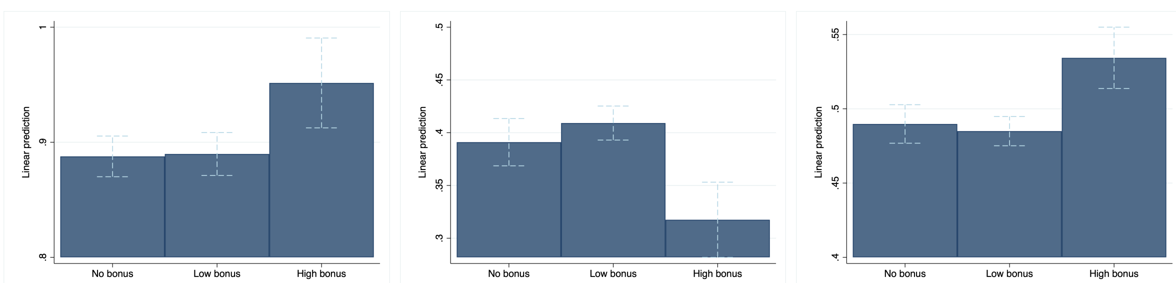

a. Vacancy filled   *b.* Teacher rank   *c.* Teacher score

Figure A.3: Contract teacher selection (Middle wage)

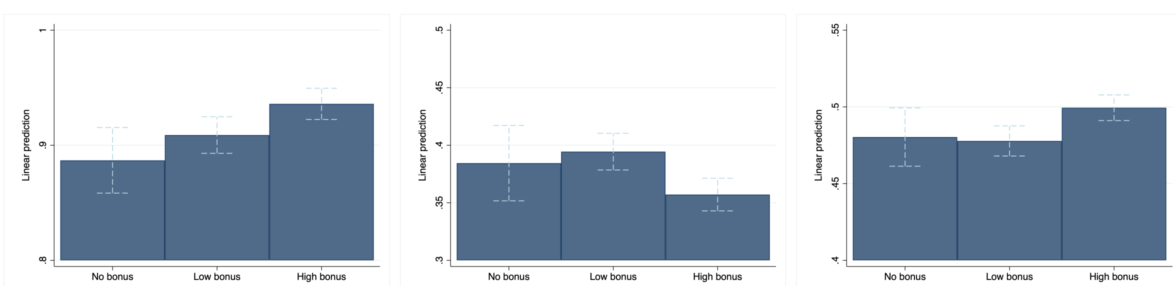

a. Vacancy filled   *b.* Teacher rank   *c.* Teacher score

Figure A.4: Contract teacher selection (High wage)

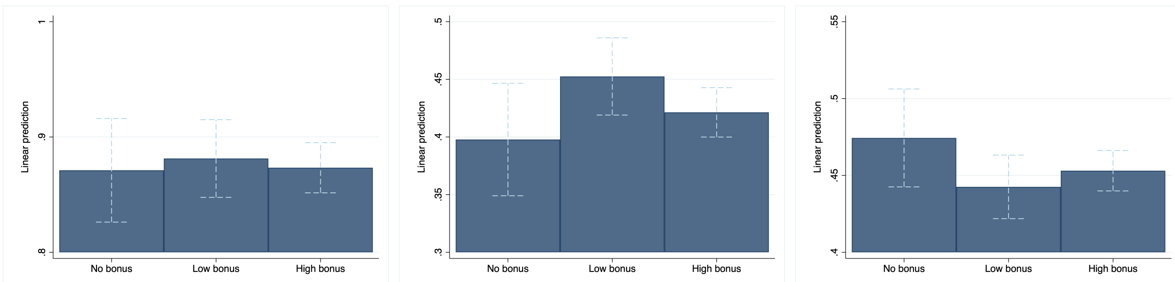

a. Vacancy filled *b.* Teacher rank *c.* Teacher score

Table A.2: Applicant Survey: Most Relevant Attributes

	All Teachers Rank				Score in Top Quartile Rank			
	1^{st}	2^{nd}	3^{rd}	In Top 3	1^{st}	2^{nd}	3^{rd}	In Top 3
Question A: Most important characteristic?								
Close to House	44.17	11.66	8.00	63.83	49.77	13.22	8.76	71.75
Safe	10.66	24.19	19.25	54.10	7.65	24.50	19.35	51.50
Well Connected	9.69	20.62	20.20	50.51	8.23	18.70	19.67	46.60
Prestige	17.92	14.12	12.29	44.33	21.13	15.77	12.68	49.58
Cultural Reasons	10.61	9.67	12.31	32.59	7.58	9.45	12.61	29.64
Good Infrastructure	2.02	8.40	12.86	23.28	1.81	7.23	11.83	20.87
Good Students	1.24	4.52	6.08	11.84	0.84	4.36	5.95	11.15
Possibility other Jobs	1.93	3.72	4.90	10.55	1.62	4.10	4.71	10.43
Career	1.76	3.10	4.09	8.95	1.36	2.67	4.44	8.47

NOTES. This table displays the share of the 5,553 survey respondents that chose the corresponding answers to Question A. The first three columns show which answer they chose and how they ranked them (by order of importance) while column 4 shows the share of respondents that listed the corresponding choice in their top 3 reasons. The last four columns display the same results for respondents that scored above the top quartile of the test score distribution for tenured teachers.

B Appendix - Additional RDD Results

Appendix Figures

Figure B.1: Manipulation charts

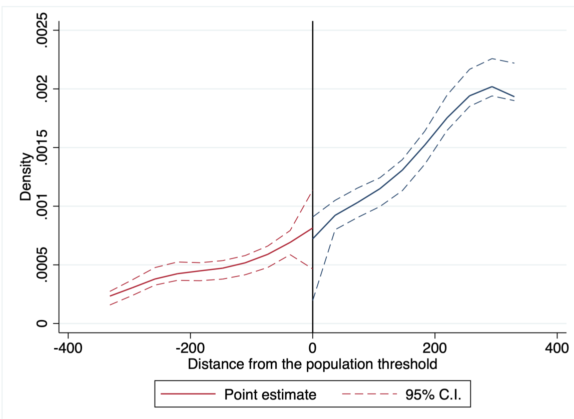

a. Population (2015)

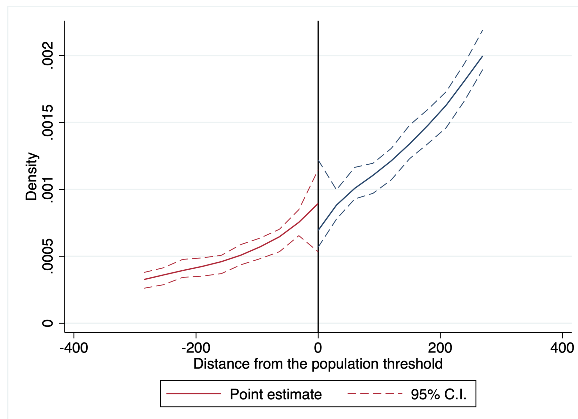

b. Population (2017)

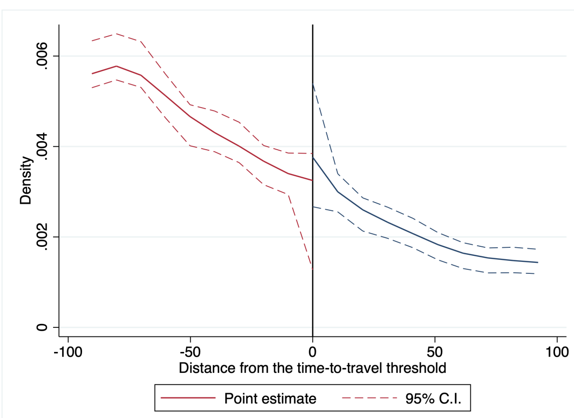

c. Time-to-travel (2015)

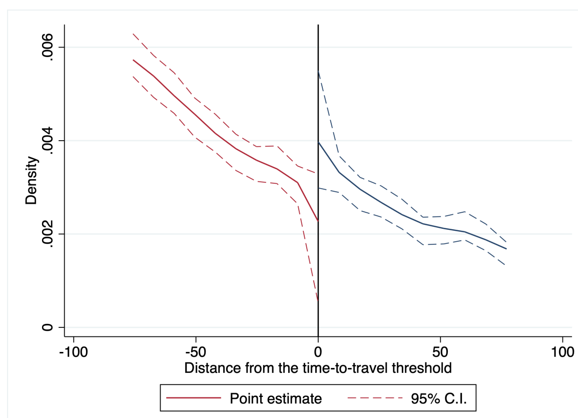

d. Time-to-travel (2017)

Notes. The figure displays the empirical densities with the corresponding confidence intervals for two running variables (population and time-to-travel) for each of the years in which the teacher recruitment drive was conducted (2015 and 2017). The density is computed using the local-polynomial estimator proposed in Cattaneo et al. (2020), and the figures show the 95% confidence intervals. The sample includes all schools with a permanent or contract teacher opening in the corresponding year.

Figure B.2: First Stage for Different Years and Treatment Status

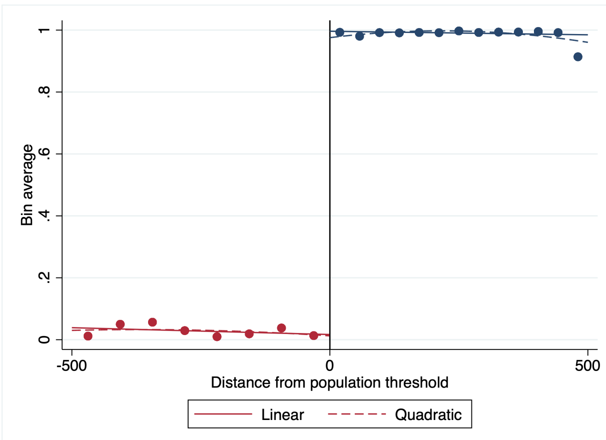

a. Treatment 2017; RV: population 2015

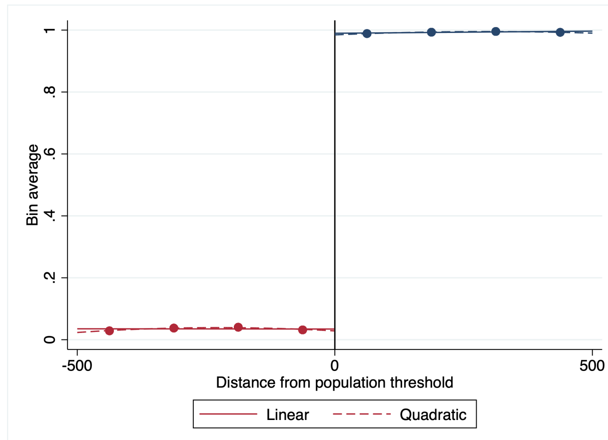

b. Treatment 2015; RV: population 2017

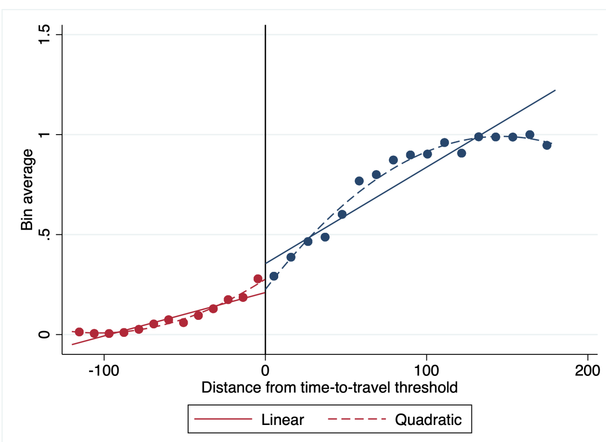

c. Treatment 2015; RV: time-to-travel 2017

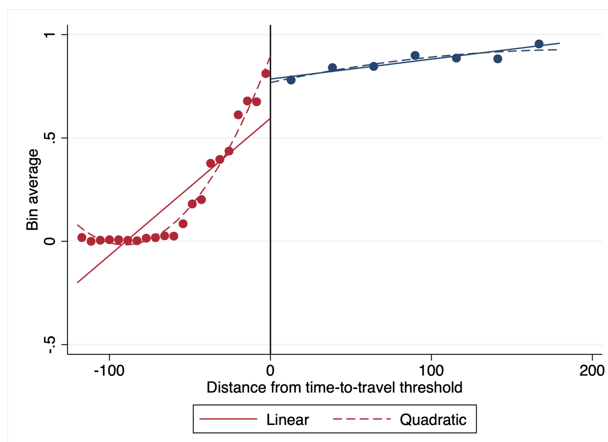

d. Treatment 2017; RV: time-to-travel 2015

Notes. The figures show the probability that a school is classified as *Extremely Rural* in each year (2015 and 2017) plotted against the two different running variables (Population and time-to-travel) for the opposite year (2017 and 2015, respectively). The regression lines are computed using linear and quadratic polynomials.

Figure B.3: Permanent Teacher Selection

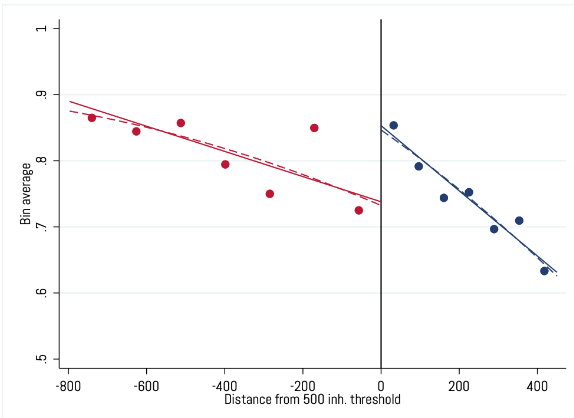

a. At least a preference

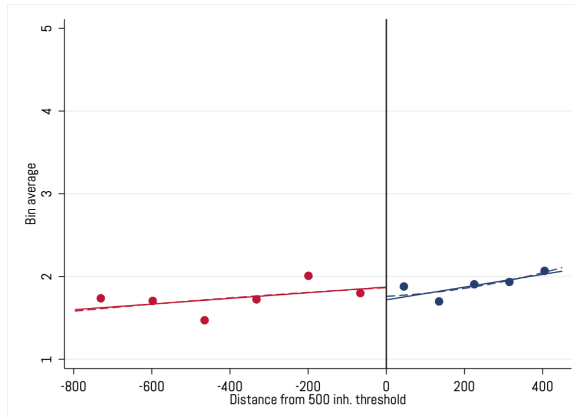

b. Highest preference

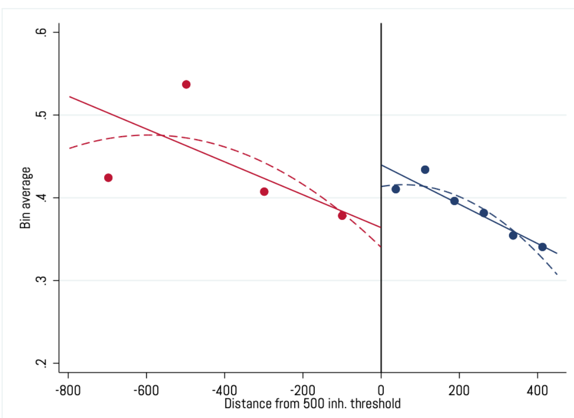

c. Vacancy filled

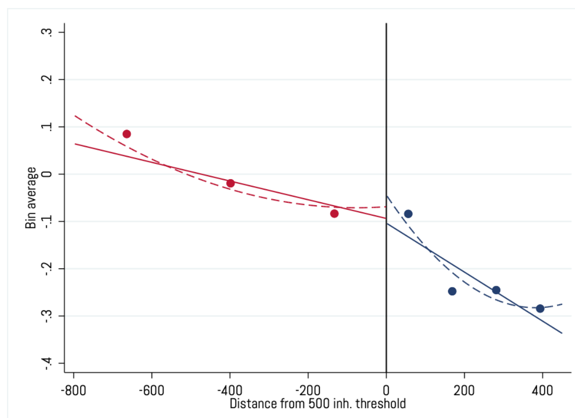

d. Teacher score

Notes. The figures show how the demand for permanent teaching positions and applicants' quality vary based on the distance from the population threshold. In subgraph (a) the outcome variable is a dummy equal to one if a school was mentioned in at least one application, while in subgraph (b) is the highest preference expressed in a teacher's list. Subgraph (c) studies whether a vacancy was filled by a certified teacher, and subgraph (d) uses as outcome variable the standardized test score obtained in the centralized test. In subgraph (b) the sample is restricted to schools that were mentioned in at least one application, while in subgraph (d) is restricted to vacancies that were actually filled by a certified teacher. Each marker indicates the average of the outcome variable within each bin, defined following the IMSE-optimal evenly spaced method by Calonico et al. (2015). The solid (dashed) lines represent the predictions from linear (quadratic) regressions estimated separately for observations to the left and to the right of the cutoff.

Figure B.4: Contract Teacher Selection

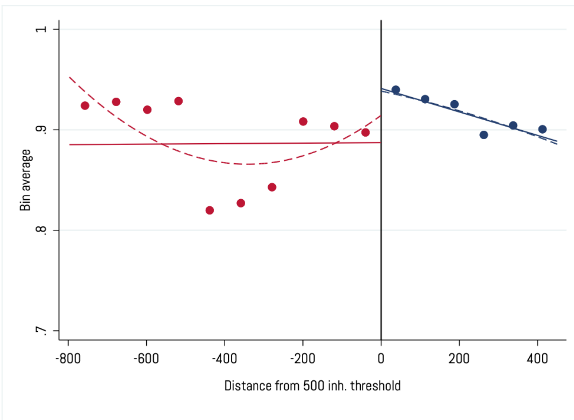

a. Vacancy filled

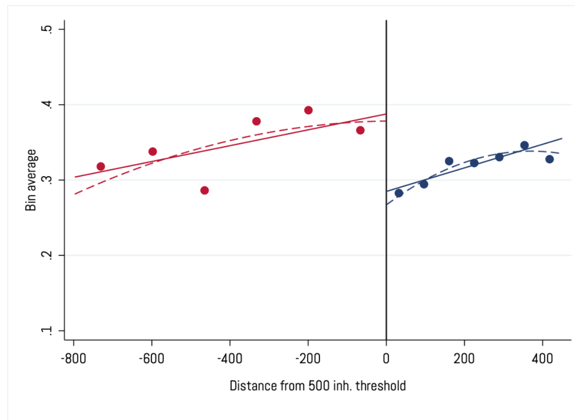

b. Teacher rank

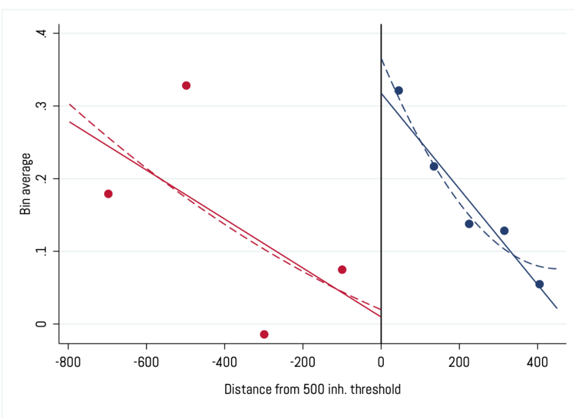

c. Teacher score

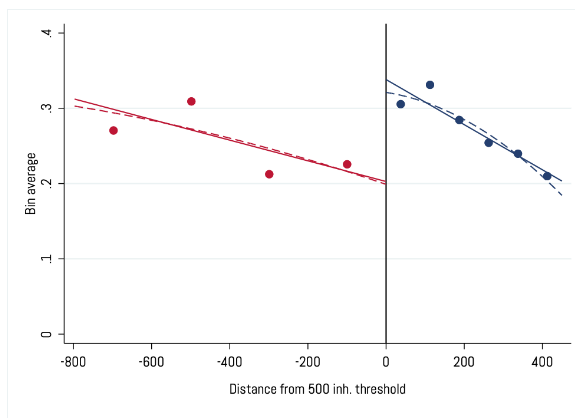

d. Top 25%

Notes. The figures show how the demand for contract teaching positions and applicants' quality vary based on the distance from the population threshold. In subgraph (a) the outcome variable is a dummy equal to one if a vacancy was filled by a certified teacher, while in subgraph (b) is the rank in which a vacancy was chosen in the deferred acceptance mechanism (normalized so that it takes value from zero to one). Subgraph (c) uses as outcome variable the standardized test score obtained in the centralized test, while subgraph (d) studies whether the teacher selected in a certain vacancy scored in the top 25% fo the distribution. In subgraphs (b) and (c) the sample is restricted to vacancies that were actually filled by a certified teacher. Each marker indicates the average of the outcome variable within each bin, defined following the IMSE-optimal evenly spaced method by Calonico et al. (2015). The solid (dashed) lines represent the predictions from linear (quadratic) regressions estimated separately for observations to the left and to the right of the cutoff.

Figure B.5: Manipulation Charts - Schools with a Contract Teacher Opening

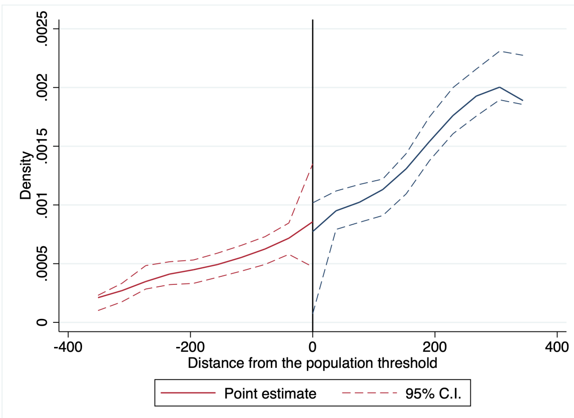

a. Population (2015)

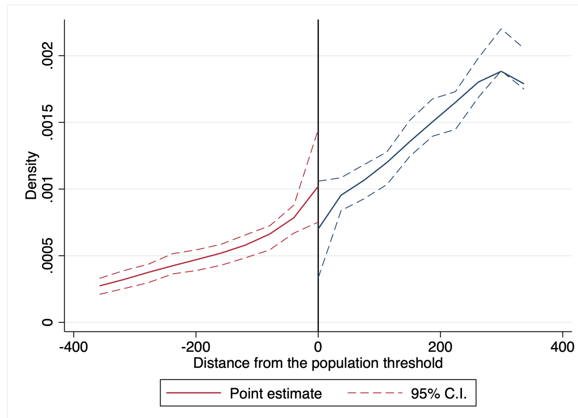

b. Population (2017)

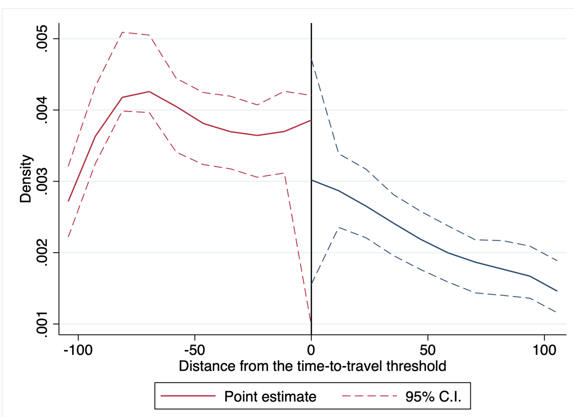

a. Time-to-travel (2015)

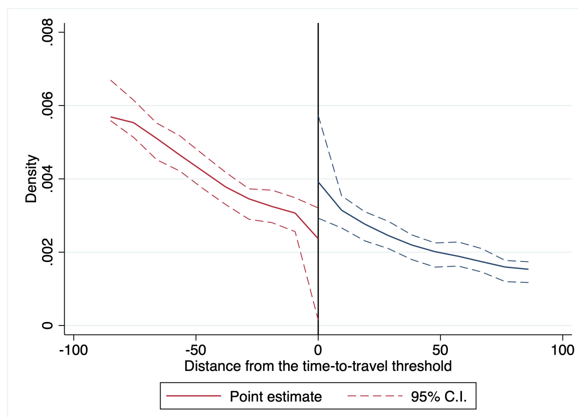

b. Time-to-travel (2017)

Notes. The figure displays the empirical densities with the corresponding confidence intervals for two running variables (population and time-to-travel) for each of the years in which the teacher recruitment drive was conducted (2015 and 2017). The density is computed using the local-polynomial estimator proposed in Cattaneo et al. (2020), and the figures show the 95% confidence intervals. The sample includes only schools with a contract teacher opening in the corresponding year.

Figure B.6: Wage Bonus and Student Achievement - Robustness to Different Specifications

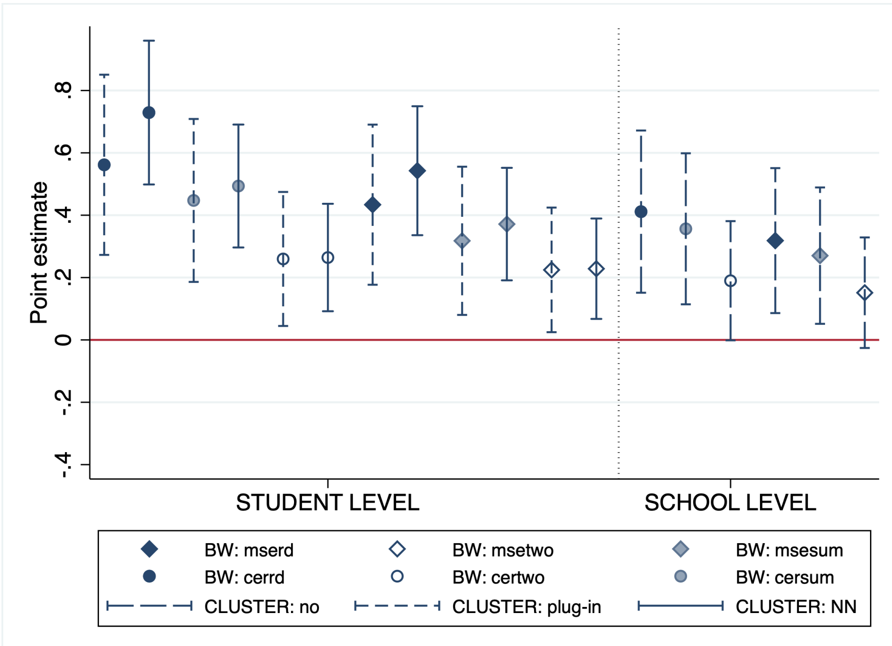

Notes. This figures shows the effect of crossing the population threshold on student achievement under different specifications. The outcome variable is the average of the standardized 2018 test scores in Math and Spanish for students in fourth grade. The sample includes schools that had an open vacancy for contract teachers the 2015 or 2017 centralized recruitment drive. Markers indicate the robust bias-corrected regression-discontinuity estimates obtained using the robust estimator proposed in Calonico et al. (2014). Regressions are defined within different specifications for the optimal bandwidths. These are: *i.* one common mean-square error (MSE) optimal bandwidth (BW: mserd); *ii.* two different MSE-optimal bandwidths, above and below the cutoff (BW: msetwo); *iii.* one common MSE-optimal bandwidth for the sum of regression estimates (BW: msesum); *iv.* one common coverage error rate (CER) optimal bandwidth (BW: cerrd); *v.* two different CER-optimal bandwidths, above and below the cutoff (BW: certwo); *vi.* one common CER-optimal bandwidth for the sum of regression estimates (BW: cersum). Vertical lines indicate confidence intervals (at the 95% level) obtained from different estimation procedures: heteroskedasticity-robust plug-in residuals (CLUSTER: no); cluster-robust plug-in residuals (CLUSTER: plug-in); cluster-robust nearest neighbor (CLUSTER: NN). The vertical dotted line separates estimates based on whether they are obtained from regressions where the unit of observation is the student (on the left) or the school (on the right). In the latter case, the outcome variables are school-level averages.

Appendix Tables

Table B.1: Covariate Smoothness around the Population Cutoff

	2015			2017		
	(1)	(2)	(3)	(4)	(5)	(6)
	Any vacancy	Permanent teacher	Contract teacher	Any vacancy	Permanent teacher	Contract teacher
Village amenities						
Electricity	0.075	0.089	0.096	0.050	0.103*	0.052
	(0.059)	(0.080)	(0.067)	(0.038)	(0.057)	(0.044)
Drinking water	0.168	0.151	0.176*	-0.013	0.043	0.015
	(0.103)	(0.151)	(0.105)	(0.066)	(0.078)	(0.075)
Sewage	0.123	0.005	0.185	0.034	-0.007	0.065
	(0.107)	(0.157)	(0.120)	(0.080)	(0.089)	(0.092)
Water tank	-0.058	-0.070	-0.085	-0.080	-0.075	-0.044
	(0.086)	(0.136)	(0.107)	(0.072)	(0.080)	(0.079)
Medical clinic	0.139	0.072	0.149	-0.037	-0.081	0.024
	(0.107)	(0.142)	(0.133)	(0.069)	(0.088)	(0.076)
Meal center	0.206**	0.163	0.164	0.116	0.119	0.134*
	(0.081)	(0.123)	(0.106)	(0.077)	(0.089)	(0.078)
Community phone	-0.026	-0.090	-0.061	0.048	0.046	0.030
	(0.049)	(0.080)	(0.050)	(0.041)	(0.053)	(0.047)
Internet access point	0.004	0.030	-0.008	0.023	0.028	0.003
	(0.080)	(0.096)	(0.089)	(0.052)	(0.068)	(0.059)
Bank	-0.023	0.031	-0.035	-0.004	0.009	0.002
	(0.037)	(0.041)	(0.048)	(0.030)	(0.034)	(0.036)
Public library	-0.035	0.009	-0.059	-0.005	-0.021	0.012
	(0.033)	(0.050)	(0.047)	(0.022)	(0.030)	(0.025)
Police	0.122	0.097	0.049	0.027	-0.014	0.021
	(0.095)	(0.146)	(0.115)	(0.065)	(0.076)	(0.079)
School amenities						
Science lab	-0.038	-0.020	-0.042	-0.052	-0.079	-0.046
	(0.069)	(0.064)	(0.083)	(0.055)	(0.059)	(0.059)
Library	-0.014	0.029	-0.137	-0.049	-0.120	-0.060
	(0.119)	(0.147)	(0.141)	(0.086)	(0.094)	(0.101)
At least a personal computer	0.102	0.147	0.135	0.069	0.147*	0.054
	(0.080)	(0.117)	(0.094)	(0.055)	(0.080)	(0.055)
Internet access	-0.125	-0.028	-0.248*	-0.078	-0.048	-0.054
	(0.116)	(0.124)	(0.138)	(0.085)	(0.098)	(0.091)
Electricity	0.173**	0.208**	0.164*	0.089*	0.124*	0.107*
	(0.082)	(0.104)	(0.093)	(0.052)	(0.064)	(0.058)
Drinking water	0.146	0.184	0.200*	0.013	0.055	0.064
	(0.102)	(0.136)	(0.109)	(0.053)	(0.082)	(0.066)
Sewage	0.036	0.036	0.041	-0.019	-0.054	0.013
	(0.090)	(0.146)	(0.114)	(0.079)	(0.084)	(0.090)
Reading room	-0.030	-0.005	-0.038	-0.016	-0.026	-0.011
	(0.041)	(0.025)	(0.051)	(0.043)	(0.049)	(0.044)
Sport pitch	0.148	0.333**	-0.026	0.062	0.008	0.055
	(0.100)	(0.149)	(0.096)	(0.084)	(0.101)	(0.087)
Courtyard	-0.016	0.151	-0.131	-0.069	-0.026	-0.113
	(0.100)	(0.127)	(0.129)	(0.079)	(0.100)	(0.086)
Gym	0.026	0.036	0.010	0.009	0.019	0.005
	(0.019)	(0.027)	(0.019)	(0.020)	(0.026)	(0.015)
Administrative office	0.046	-0.020	0.060	0.085	0.047	0.086
	(0.102)	(0.146)	(0.107)	(0.078)	(0.099)	(0.091)
Pool	0.041	0.047	0.016	-0.050	-0.086*	-0.021
	(0.048)	(0.036)	(0.058)	(0.047)	(0.052)	(0.054)
Courtyard	-0.062	-0.058	-0.062	-0.073**	-0.039	-0.093**
	(0.038)	(0.058)	(0.054)	(0.032)	(0.034)	(0.043)
Resting room	0.019	0.036	0.006	0.120**	0.074	0.135**
	(0.064)	(0.082)	(0.068)	(0.057)	(0.066)	(0.060)
Breastfeeding room	-0.005	0.046	-0.099	-0.044	-0.091	-0.033
	(0.086)	(0.110)	(0.102)	(0.070)	(0.077)	(0.075)
Courtyard	0.035	0.025	0.026	0.000	0.023	-0.002
	(0.025)	(0.033)	(0.030)	(0.024)	(0.025)	(0.026)
Dining hall	0.035	0.229*	-0.075	-0.043	-0.107	-0.056
	(0.095)	(0.137)	(0.110)	(0.083)	(0.088)	(0.095)

Table B.2: Probability of Openings around the Population Cutoff

	All		Permanent teacher		Contract teacher	
	(1)	(2)	(3)	(4)	(5)	(6)
	Vacancy	N. of vacancies	Vacancy	N. of vacancies	Vacancy	N. of vacancies
Above cutoff	-0.007	-0.114	0.008	-0.042	-0.007	-0.113
	(0.040)	(0.138)	(0.041)	(0.091)	(0.044)	(0.135)
Mean dep. var. (LHS)	0.476	0.960	0.252	0.463	0.397	0.764
BW	245.255	185.331	166.599	172.265	221.811	184.597
Observations	6196	4244	3793	3904	5365	4221

Notes. This table reports the effect of crossing the population threshold on the probability that vacancy is posted (and their number) in the 2015 or 2017 assignment process. In column (1) the outcome variable is a dummy equal to 1 if the school had at least a vacancy (of any type), while in column (2) is the number of open vacancies. Columns (3)-(4) and (5)-(6) are the analogous of columns (1)-(2) but focus only on permanent and contract teachers vacancies, respectively. Cells report the bias-corrected regression-discontinuity estimates obtained using the robust estimator proposed in Calonico et al. (2014). Regressions are defined within a mean-square error optimal bandwidth (BW), reported at the bottom part of the table. The table also reports the mean of the dependent variable computed within the intervals $(0, +BW)$ (right-hand-side of the cutoff) and $(-BW, 0]$ (left-hand-side of the cutoff). SE are clustered at the school level. *** p< 0.01, ** p<0.05, and *p<0.10.

Table B.3: Monetary Incentives and Teacher Selection (2015)

Panel A: Permanent teacher				
	(1)	(2)	(3)	(4)
	At least a preference	Highest preference	Vacancy filled	Teacher score (std)
Above cutoff	0.071	0.191	-0.182	0.419
	(0.083)	(0.334)	(0.122)	(0.414)
Bounds	[.085; .194]	[-.211; .26]	[-.149; -.07]	[.266; .266]
Mean dep. var. (LHS)	0.795	1.556	0.508	0.277
BW	247.770	150.323	219.949	134.609
Schools	590	301	488	268
Observations	590	248	661	189
	Panel B: Contract teacher			
	(1)	(2)	(3)	(4)
	Vacancy filled	Teacher rank	Teacher score (std)	Top 25%
Above cutoff	0.098	-0.132**	0.644***	0.200**
	(0.067)	(0.060)	(0.196)	(0.082)
Bounds	[.088; .088]	[-.166; -.097]	[.482; .722]	[.187; .187]
Mean dep. var. (LHS)	0.870	0.391	-0.104	0.142
BW	202.630	170.661	150.932	145.821
Schools	593	481	424	399
Observations	987	720	651	688

Notes. This table reports the effect of crossing the population threshold on different outcomes of the 2015 centralized recruitment drive. Panel A uses the sample of permanent teachers. In column (1) the outcome variable is a dummy equal to one if a school was mentioned in at least one application, while column (2) uses the highest preference expressed in a teacher's list. Column (3) studies whether a vacancy was filled by a certified teacher, and Column (4) uses as outcome variable the standardized test score obtained in the centralized test. In Column (2) the sample is restricted to schools that were mentioned in at least one application, while in Column (4) is restricted to vacancies that were actually filled by a certified teacher. Panel B focuses on the selection process of contract teachers. Columns (1) and (3) are analogous to Columns (3) and (4) from panel A. Column (2) shows the effects on the rank in which a vacancy was chosen in the deferred acceptance mechanism (normalized so that it takes value from zero to one). Column (4) is an indicator for whether the teacher selected in a certain vacancy scored in the top 25% fo the distribution. Cells report the bias-corrected regression-discontinuity estimates obtained using the robust estimator proposed in Calonico et al. (2014) and their bounds estimated using the procedure developed in (Gerard et al. 2020). Regressions are defined within a mean-square error optimal bandwidth (BW), reported at the bottom part of the table. The table also reports the mean of the dependent variable computed within the intervals $(0, +BW)$ (right-hand-side of the cutoff) and $(-BW, 0]$ (left-hand-side of the cutoff). SE are clustered at the school level. *** p< 0.01, ** p<0.05, and *p<0.10.

Table B.4: Monetary Incentives and Teacher Selection (2017)

	(1)	(2)	(3)	(4)
Panel A: Permanent teacher				
	At least a preference	Highest preference	Vacancy filled	Teacher score (std)
Above cutoff	0.252***	-0.087	0.069	-0.085
	(0.087)	(0.309)	(0.082)	(0.209)
Bounds	[.199; .382]	[-.882; .399]	[.035; .113]	[-.508; .402]
Mean dep. var. (LHS)	0.743	1.930	0.330	-0.177
BW	157.073	147.567	170.460	164.311
Schools	626	585	681	662
Observations	626	455	1261	456
	Panel B: Contract teacher			
	(1)	(2)	(3)	(4)
	Vacancy filled	Teacher rank	Teacher score (std)	Top 25%
Above cutoff	0.020	-0.120***	0.375**	0.091
	(0.056)	(0.042)	(0.149)	(0.066)
Bounds	[.021; .021]	[-.112; -.112]	[.365; .365]	[.106; .106]
Mean dep. var. (LHS)	0.912	0.360	0.166	0.264
BW	157.377	167.063	185.715	184.230
Schools	800	855	941	935
Observations	1435	1416	1548	1648

Notes. This table reports the effect of crossing the population threshold on different outcomes of the 2017 centralized recruitment drive. Panel A uses the sample of permanent teachers. In column (1) the outcome variable is a dummy equal to one if a school was mentioned in at least one application, while column (2) uses the highest preference expressed in a teacher's list. Column (3) studies whether a vacancy was filled by a certified teacher, and Column (4) uses as outcome variable the standardized test score obtained in the centralized test. In Column (2) the sample is restricted to schools that were mentioned in at least one application, while in Column (4) is restricted to vacancies that were actually filled by a certified teacher. Panel B focuses on the selection process of contract teachers. Columns (1) and (3) are analogous to Columns (3) and (4) from panel A. Column (2) shows the effects on the rank in which a vacancy was chosen in the deferred acceptance mechanism (normalized so that it takes value from zero to one). Column (4) is an indicator for whether the teacher selected in a certain vacancy scored in the top 25% fo the distribution. Cells report the bias-corrected regression-discontinuity estimates obtained using the robust estimator proposed in Calonico et al. (2014) and their bounds estimated using the procedure developed in (Gerard et al. 2020). Regressions are defined within a mean-square error optimal bandwidth (BW), reported at the bottom part of the table. The table also reports the mean of the dependent variable computed within the intervals $(0, +BW)$ (right-hand-side of the cutoff) and $(-BW, 0]$ (left-hand-side of the cutoff). SE are clustered at the school level. *** p< 0.01, ** p<0.05, and *p<0.10.

Table B.5: Monetary Incentives and the Selection of Permanent Teachers - low *vs* high bonus

Panel A: Low bonus				
	(1)	(2)	(3)	(4)
	At least a preference	Highest preference	Vacancy filled	Teacher score (std)
Above cutoff	0.205***	0.489*	0.016	0.397
	(0.073)	(0.283)	(0.115)	(0.255)
Bounds	[.183; .257]	[-.234; .843]	[.002; .002]	[.219; .305]
Mean dep. var. (LHS)	0.821	1.699	0.463	-0.025
BW	228.808	155.328	172.984	207.827
Schools	548	378	406	488
Observations	620	367	672	418
	Panel B: High bonus			
	(1)	(2)	(3)	(4)
	At least a preference	Highest preference	Vacancy filled	Teacher score (std)
Above cutoff	0.150	-0.442	-0.020	-0.480**
	(0.098)	(0.347)	(0.089)	(0.234)
Bounds	[.144; .214]	[-.966; -.094]	[-.005; .055]	[-.654; -.059]
Mean dep. var. (LHS)	0.706	2.000	0.272	-0.135
BW	189.312	161.292	130.572	152.920
Schools	545	446	344	418
Observations	666	396	779	310

Notes. This table reports the effect of crossing the population threshold on different outcomes, separately for schools below (+S/ 30 bonus, in Panel A) and above (+S/ 400 bonus, in Panel A) the time-to-travel threshold. In both panels, the sample includes all permanent teacher vacancies. In column (1) the outcome variable is a dummy equal to one if a school was mentioned in at least one application, while column (2) uses the highest preference expressed in a teacher's list. Column (3) studies whether a vacancy was filled by a certified teacher, and Column (4) uses as outcome variable the standardized test score obtained in the centralized test. In Column (2) the sample is restricted to schools that were mentioned in at least one application, while in Column (4) is restricted to vacancies that were actually filled by a certified teacher. Cells report the bias-corrected regression-discontinuity estimates obtained using the robust estimator proposed in Calonico et al. (2014) and their bounds estimated using the procedure developed in (Gerard et al. 2020). Regressions are defined within a mean-square error optimal bandwidth (BW), reported at the bottom part of the table. The table also reports the mean of the dependent variable computed within the intervals $(0, +BW)$ (right-hand-side of the cutoff) and $(-BW, 0]$ (left-hand-side of the cutoff). SE are clustered at the school×year level. *** p< 0.01, ** p<0.05, and *p<0.10.

Table B.6: Monetary Incentives and the Selection of Contract Teachers - low *vs* high bonus

Panel A: Low bonus				
	(1)	(2)	(3)	(4)
	Vacancy filled	Teacher rank	Teacher score (std)	Top 25%
Above cutoff	-0.061	-0.147***	0.462**	0.081
	(0.044)	(0.050)	(0.184)	(0.090)
Bounds	[-.059; -.059]	[-.145; -.145]	[.493; .493]	[.098; .098]
Mean dep. var. (LHS)	0.939	0.328	0.266	0.306
BW	172.714	107.926	122.199	153.501
Schools	581	353	394	512
Observations	1008	608	671	904
	Panel B: High bonus			
	(1)	(2)	(3)	(4)
	Vacancy filled	Teacher rank	Teacher score (std)	Top 25%
Above cutoff	0.130**	-0.146***	0.593***	0.176***
	(0.062)	(0.045)	(0.172)	(0.066)
Bounds	[.134; .18]	[-.198; -.106]	[.317; .813]	[.138; .187]
Mean dep. var. (LHS)	0.862	0.420	-0.127	0.142
BW	162.363	215.588	159.784	160.932
Schools	471	654	457	461
Observations	1300	1610	1150	1275

Notes. This table reports the effect of crossing the population threshold on different outcomes, separately for schools below (+S/ 30 bonus, in Panel A) and above (+S/ 400 bonus, in Panel A) the time-to-travel threshold. In both panels, the sample includes all contract teacher vacancies. In column (1) the outcome variable is a dummy equal to one if a vacancy was filled by a certified teacher. Column (2) shows the effects on the rank in which a vacancy was chosen in the deferred acceptance mechanism (normalized so that it takes value from zero to one), and Column (3) uses as outcome variable the standardized test score obtained in the centralized test. Column (4) is an indicator for whether the teacher selected in a certain vacancy scored in the top 25% fo the distribution. In Columns (2) and (3) the sample is restricted to vacancies that were actually filled by a certified teacher. Cells report the bias-corrected regression-discontinuity estimates obtained using the robust estimator proposed in Calonico et al. (2014) and their bounds estimated using the procedure developed in (Gerard et al. 2020). Regressions are defined within a mean-square error optimal bandwidth (BW), reported at the bottom part of the table. The table also reports the mean of the dependent variable computed within the intervals $(0, +BW)$ (right-hand-side of the cutoff) and $(-BW, 0]$ (left-hand-side of the cutoff). SE are clustered at the school×year level. *** p< 0.01, ** p<0.05, and *p<0.10.

Table B.7: Monetary Incentives and Teaching Staff Composition

	Permanent teacher			Contract teacher		
	(1)	(2)	(3)	(4)	(5)	(6)
	N. of teachers	Student/Teacher	% of permanent t.	N. of teachers	Student/Teacher	% of contract t.
Above cutoff	0.189	-0.111	0.082*	-0.541	0.048	-0.039
	(0.352)	(0.184)	(0.043)	(0.373)	(0.184)	(0.038)
Mean dep. var. (LHS)	6.617	2.667	0.543	6.548	2.598	0.409
BW	177.285	145.189	243.788	145.863	168.266	182.564
Observations	1068	841	1648	1120	1304	1441

Notes. This table reports the effect of crossing the population threshold on the number and the composition of teaching staff in schools that had an open vacancy in the 2015 or 2017 assignment process. The sample in columns (1) to (3) includes schools that had vacancies for permanent teachers. In column (1) the outcome variable is the total number of teachers, in column (2) is the students to teachers ratio, while in column (3) is the share of permanent teachers. Columns (4) to (6) are the analogous of columns (1)-(3) for schools that had vacancies for contract teachers. Cells report the bias-corrected regression-discontinuity estimates obtained using the robust estimator proposed in Calonico et al. (2014). Regressions are defined within a mean-square error optimal bandwidth (BW), reported at the bottom part of the table. The table also reports the mean of the dependent variable computed within the intervals $(0, +BW)$ (right-hand-side of the cutoff) and $(-BW, 0]$ (left-hand-side of the cutoff). SE are clustered at the school level. *** p< 0.01, ** p<0.05, and *p<0.10.

Table B.8: Monetary Incentives and the Selection of Contract Teachers - Other Teachers' Characteristics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
						Years as public school teacher		
	Female	Age	Indigenous	University	Novice Teacher	0	1-3	> 3
Above cutoff	0.103*	-1.526*	-0.017	0.145	0.013	0.052	0.069	-0.135**
	(0.059)	(0.860)	(0.127)	(0.101)	(0.037)	(0.038)	(0.048)	(0.062)
Mean dep. var. (LHS)	0.585	38.851	0.363	0.230	0.045	0.153	0.391	0.388
BW	131.823	174.296	192.388	142.891	194.678	142.026	234.147	140.239
Schools	757	1025	1149	818	1169	818	1446	805
Observations	1769	2297	847	613	853	1922	3281	1892

NOTES. This table reports the effect of crossing the population threshold on several teachers' characteristics. These are a female dummy (column 1), age (column 2), a dummy equal to 1 if the teacher speaks a peruvian indigenous language (column 3), an indicator for university or technical institute education (column 4), and a dummy equal to 1 if the teacher has no previous teaching experience, neither in the public nor private sector (column 5). In columns 6 to 8, the outcome variables are a set of binary indicators the number of years of teaching experience in the public sector, measured as the number of years the teacher was observed in the teacher occupation and payroll system (*NEXUS*) before the assignment process.The sample includes all contract teacher vacancies assigned in the 2015 and 2017 processes, regardless of whether they were assigned to a certified or non-certified teachers. In columns (3), (4), and (5) the sample includes only vacancies assigned to certified teachers in 2015, as the same information is not available for the 2017 assignment process. Cells report the bias-corrected regression-discontinuity estimates obtained using the robust estimator proposed in Calonico et al. (2014). Regressions are defined within a mean-square error optimal bandwidth (BW), reported at the bottom part of the table. The table also reports the mean of the dependent variable computed within the intervals $(0, +BW)$ (right-hand-side of the cutoff) and $(-BW, 0]$ (left-hand-side of the cutoff). SE are clustered at the school×year level. *** p< 0.01, ** p<0.05, and *p<0.10.

Table B.9: Origin of Newly Recruited Teachers: Regression Estimates

	Unfilled v.	Filled vacancy											
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
		Novice t.	Same sch.	0-99	100-199	200-299	300-399	400-499	500-599	600-699	700-799	800-2000	Urban
Above cutoff	-0.043	0.049	-0.002	-0.023	-0.009	0.037	0.032	0.006	0.003	-0.027	-0.035*	0.026	-0.008
	(0.045)	(0.041)	(0.039)	(0.027)	(0.035)	(0.032)	(0.026)	(0.019)	(0.022)	(0.018)	(0.018)	(0.019)	(0.028)
Mean dep. var. (LHS)	0.102	0.216	0.154	0.051	0.075	0.087	0.047	0.044	0.049	0.025	0.024	0.056	0.066
BW	160.098	122.583	163.318	144.407	129.042	154.888	124.410	211.734	173.911	162.758	123.568	155.949	143.211
Schools	943	693	969	826	742	905	711	1281	1018	961	700	911	822
Observations	2218	1692	2278	1975	1795	2146	1729	2962	2366	2263	1708	2154	1966

Notes. This table reports the effect of crossing the population threshold on a set of indicators for the teachers' location in the year before the assignment process. These are a dummy equal to one if the vacancy is filled by a teacher already in the same school (column 3), or is filled by a teacher whose previous location falls into the population bin indicated in the column header (columns 4-13). Urban schools are those in localities above 2000 inhabitants. The table also reports the effect of crossing the population threshold on the probability that the vacancy remains unfilled (column 1), or is filled by a new entrant in the public education system (column 2). Teachers' previous school is determined based on the teacher occupation and payroll system (*NEXUS*). The sample includes all contract teacher vacancies assigned to a certified teacher in the 2015 and 2017 processes. Cells report the bias-corrected regression-discontinuity estimates obtained using the robust estimator proposed in Calonico et al. (2014). Regressions are defined within a mean-square error optimal bandwidth (BW), reported at the bottom part of the table. The table also reports the mean of the dependent variable computed within the intervals $(0, +BW)$ (right-hand-side of the cutoff) and $(-BW, 0]$ (left-hand-side of the cutoff). SE are clustered at the school×year level. *** p< 0.01, ** p<0.05, and *p<0.10.

Table B.10: Wage Bonus and Student Achievement

	No vacancy	Vacancy		
Panel A: Spanish				
	(1)	(2)	(3)	(4)
		Any vacancy	Permanent teacher	Contract teacher
Above cutoff	0.014	0.298**	-0.057	0.317**
	(0.157)	(0.127)	(0.190)	(0.137)
Mean dep. var. (LHS)	-0.471	-0.470	-0.382	-0.491
BW	124.095	108.453	175.257	114.883
Schools	372	691	292	622
Observations	3948	9700	3409	8966
Panel B: Math				
	No vacancy	Vacancy		
	(1)	(2)	(3)	(4)
		Any vacancy	Permanent	Contract
Above cutoff	0.039	0.350**	-0.047	0.470***
	(0.174)	(0.142)	(0.248)	(0.159)
Mean dep. var. (LHS)	-0.416	-0.416	-0.296	-0.417
BW	126.632	112.196	163.972	101.761
Schools	381	710	275	561
Observations	4046	10013	3205	8146

Notes. This table reports the effect of crossing the population threshold on student achievement in Math and Spanish. In all columns, the outcome variable is the standardized 2018 test scores in Spanish (Panel A) and Math (Panel B) for students in fourth grade. The sample in Columns and (2) is split based on whether the school had an open vacancy (of any type) in the 2015 or 2017 centralized recruitment driver. In Column (3) and (4), the sample is further restricted to schools that had vacancies for permanent or contract teachers, respectively. Each cell reports the bias-corrected regression-discontinuity estimates obtained using the robust estimator proposed in Calonico et al. (2014). Regressions are defined within a mean-square error optimal bandwidth (BW), reported at the bottom part of the table. The table also reports the mean of the dependent variable computed within the intervals $(0, +BW)$ (right-hand-side of the cutoff) and $(-BW, 0]$ (left-hand-side of the cutoff). SE are clustered at the school×year level. *** p< 0.01, ** p<0.05, and *p<0.10.

Table B.11: Wage Bonus and Students' Achievement Level

	(1)	(2)	(3)	(4)
Panel A: Spanish				
	Below basic	Basic	Intermediate	Proficient
Above cutoff	-0.131**	0.009	0.074**	0.058
	(0.055)	(0.033)	(0.034)	(0.036)
Mean dep. var. (LHS)	0.247	0.341	0.241	0.178
BW	135.189	150.929	141.445	117.246
Schools	742	840	776	638
Observations	10429	11632	10815	9134
Panel B: Math				
	(1)	(2)	(3)	(4)
	Below basic	Basic	Intermediate	Proficient
Above cutoff	-0.117**	-0.053	0.080**	0.119**
	(0.055)	(0.036)	(0.039)	(0.050)
Mean dep. var. (LHS)	0.204	0.282	0.345	0.181
BW	132.558	124.554	152.884	99.264
Schools	730	681	855	549
Observations	10270	9702	11849	7985

Notes. This table reports the effect of crossing the population threshold on student achievement in Math and Spanish classified according to four categories. These are below basic (*Previo al inicio*), basic (*En inicio*), intermediate (*En proceso*), and proficient (*Satisfactorio*). In each column, the outcome variable is a dummy equal to one if a (fourth-grade) student falls into the corresponding category. The sample includes schools with an open position for contract teachers. Each cell reports the bias-corrected regression-discontinuity estimates obtained using the robust estimator proposed in Calonico et al. (2014). Regressions are defined within a mean-square error optimal bandwidth (BW), reported at the bottom part of the table. The table also reports the mean of the dependent variable computed within the intervals $(0, +BW)$ (right-hand-side of the cutoff) and $(-BW, 0]$ (left-hand-side of the cutoff). SE are clustered at the school×year level. *** p< 0.01, ** p<0.05, and *p<0.10.

C Appendix - Additional Model Results

Demand Model Parameter Estimates

Table C.1: Preference Estimates

		Heterogeneity			
	All	Base	Age < 30	Rural	Male
Wage	0.635***	0.111	0.243	0.134	0.983***
	(0.0657)	(0.105)	(0.205)	(0.146)	(0.137)
Poverty Score Locality	-0.0817***	-0.108***	-0.0201	0.0704***	0.0537***
	(0.00750)	(0.0106)	(0.0224)	(0.0182)	(0.0162)
School Infrastructure Index	0.0177	0.0376**	0.0208	-0.00236	-0.0514*
	(0.00975)	(0.0137)	(0.0281)	(0.0233)	(0.0209)
Locality Infrastructure Index	0.0510***	0.0397	0.0520	0.0128	-0.0225
	(0.0148)	(0.0212)	(0.0440)	(0.0350)	(0.0313)
Spline: Distance to place of origin					
$. < 20$km	-141.2***	-141.4***	-2.473	-31.98***	23.89***
	(1.573)	(2.185)	(4.480)	(3.682)	(3.465)
20km $\leq . <$ 100km	-34.69***	-30.52***	2.616*	-10.65***	-3.114**
	(0.458)	(0.667)	(1.326)	(1.102)	(0.965)
100km $\leq . <$ 200km	-15.36***	-16.03***	-2.022	3.574**	-0.699
	(0.512)	(0.729)	(1.508)	(1.299)	(1.068)
200km $\leq . <$ 300km	-16.14***	-18.39***	5.170**	0.384	4.372**
	(0.663)	(0.951)	(1.891)	(1.689)	(1.372)
$. \geq$ 300km	-2.612***	-1.839***	-0.932*	-0.731	-1.416***
	(0.142)	(0.186)	(0.425)	(0.403)	(0.309)
Match effects					
Mestizo × Share Mestizo in school's district	0.460***	0.296***	-0.393*	0.405**	0.232
	(0.0619)	(0.0877)	(0.176)	(0.154)	(0.131)
Quechua × Share Quechua in school's district	1.670***	2.416***	-0.125	-1.538***	-0.306
	(0.105)	(0.173)	(0.313)	(0.217)	(0.212)
× Bilingual school	1.106***	1.269***	0.0686	-0.159	-0.291*
	(0.0654)	(0.103)	(0.200)	(0.144)	(0.136)
Aimara × Share Aimara in school's district	2.338***	3.560***	-2.547***	-0.159	-0.649
	(0.280)	(0.449)	(0.632)	(0.611)	(0.550)
× Bilingual school	0.894***	0.959***	0.799	-0.214	-0.202
	(0.167)	(0.259)	(0.460)	(0.372)	(0.340)
Wage bonus determinants					
log(Population)	0.0360	0.0206	-0.0580	-0.0496	0.152***
	(0.0206)	(0.0291)	(0.0606)	(0.0549)	(0.0454)
log(Population)^2	-0.0139***	-0.0124***	0.00526	0.00678*	-0.0107***
	(0.00100)	(0.00137)	(0.00296)	(0.00302)	(0.00226)
Time to closest city	-0.0575***	-0.0809***	-0.120**	0.0385	0.0362
	(0.00949)	(0.0162)	(0.0367)	(0.0226)	(0.0200)
Time to closest city^2	-0.00118***	-0.00124***	-0.000554	0.000321	-0.0000381
	(0.0000829)	(0.000160)	(0.00525)	(0.000237)	(0.000196)
log(Pop) × Time	0.0261***	0.0296***	0.0194***	-0.0131***	-0.00365
	(0.00144)	(0.00233)	(0.00525)	(0.00345)	(0.00297)
VRAEM	0.316***	0.455***	-0.155	0.0445	-0.266**
	(0.0472)	(0.0713)	(0.157)	(0.111)	(0.0978)
Frontera	-0.0472	-0.132*	-0.0530	0.220*	0.105
	(0.0439)	(0.0670)	(0.126)	(0.105)	(0.0923)
Multigrado	-0.519***	-0.535***	-0.000266	0.0879	-0.000256
	(0.0287)	(0.0439)	(0.0851)	(0.0630)	(0.0594)
Unidocente	-0.875***	-1.097***	-0.0117	0.142	0.381***
	(0.0451)	(0.0730)	(0.142)	(0.0989)	(0.0939)
Bilingual school	-0.762***	-0.961***	-0.359**	0.477***	0.132
	(0.0352)	(0.0560)	(0.110)	(0.0786)	(0.0738)

NOTES. Standard errors in parenthesis. *** $p < 0.001$, ** $p < 0.01$, and * $p < 0.05$. This table displays the estimates of the model for teachers' preferences described in Section 4.1. We use the 8,190 teachers assigned in 2015 along with the 10,569 teachers assigned in 2017 and consider the two samples as independent cross sections. We construct the feasible choice sets by first determining the score of the lowest ranked applicant in each school (which we will call cutoffs). If the school hasn't filled all vacancies it is feasible by definition, if the school is full, it is feasible only if the teacher has a score above the cutoff. We then estimate this discrete choice model with personalized choice sets by maximum likelihood. Across both years, the feasible choice sets contain 5,672 schools on average.

Additional Results from Counterfactuals and Simulations

Figure C.1: Distribution of Wage Bonuses Under Counterfactual Policy

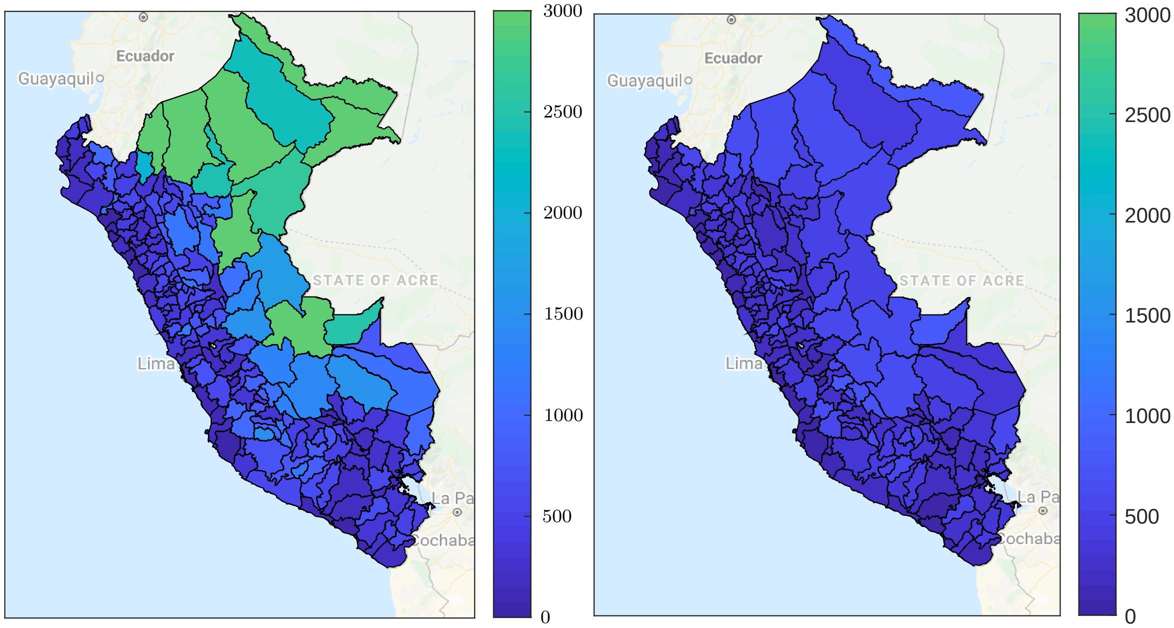

Policy Equalizing Teacher Quality **Current Policy**

NOTES. The first panel shows the monthly wage bonuses (in soles) needed to fill every vacancy in the 2015 and the 2017 concurso with an above median teacher averaged at the province level. The second panel maps the monthly wage bonuses offered by the current policy averaged at the province level.