

Teacher Compensation and Structural Inequality: Evidence from Centralized Teacher School Choice in Perú*

Matteo Bobba[†] Tim Ederer[‡] Gianmarco León-Ciliotta[§]
Christopher A. Neilson[¶] Marco Nieddu^{||}

August 7, 2026

Abstract

We exploit data on centralized teacher recruitment in Perú to establish that wage rigidity creates large urban-rural disparities in teacher effectiveness. Leveraging a teacher compensation reform, we provide causal evidence that increasing salaries in less desirable locations attracts qualified teachers and improves student learning. We estimate a model of teacher sorting and student achievement featuring rich heterogeneity in teachers' preferences and effectiveness. Substantial equity and efficiency gains arise from reallocating existing teachers to exploit match effects or attracting applicants with higher average effectiveness into public teaching. Counterfactual wage schedules achieving these gains suggest that the latter is more cost-effective.

*We thank the Ministry of Education of Perú, and in particular Camila Alva, for supporting rigorous policy evaluation efforts and continued collaboration. Mohammad Akbarpour, Ernesto Dal Bo, Michael Dinerstein, Pascaline Dupas, Federico Finan, Yinghua He, Imran Rasul, Jesse Silbert, Alessandro Tarozzi, and Takuro Yamashita generously provided us with insightful feedback on earlier versions of this draft. We also thank Mariel Bedoya, Manuela Cardona Badillo, Maria Elena Guerrero, Isabel Jacas, Cecilia Moreira, and Eduardo Vargas for outstanding research assistance. Matteo Bobba acknowledges 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). Gianmarco 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 CEX2019-000915-S), grant PID2023-150768NB-I00 and RyC-2017-23172 financed by MICIU/AEI /10.13039/501100011033 and FSE invierte en tu futuro, and the BSE Seed Grant.

[†]Toulouse School of Economics & CEPR, matteo.bobba@tse-fr.eu

[‡]Carnegie Mellon University, tederer@andrew.cmu.edu

[§]Universitat Pompeu Fabra & BSE & IPEG & BREAD & CEPR, gianmarco.leon@upf.edu

[¶]Yale University & NBER, christopher.neilson@yale.edu

^{||}University of Cagliari & CRENoS, mgnieddu@unica.it

1 Introduction

Children born in remote or rural communities attain significantly lower academic achievement than their urban counterparts (World Bank, 2018). These disparities are particularly salient in developing countries, where remote locations suffer from a structural and historically persistent underdevelopment (Sokoloff and Engerman, 2000; Banerjee and Iyer, 2005; Dell, 2010). Designing effective policies that provide equal opportunities despite these barriers is thus a first-order concern. In this paper, we show that structural inequalities in schooling outcomes can be mitigated or reinforced by the inherent mobility of a critical factor in the production of human capital: teachers.

Teachers are a key determinant of student learning and long-run outcomes (Rivkin et al., 2005; Chetty et al., 2014b; Araujo et al., 2016; Jackson, 2018). Recent evidence further documents that teachers hold comparative advantages in teaching different types of students (Gershenson et al., 2022; Ahn et al., 2025; Graham et al., 2023), implying that teacher allocation is not only relevant for equity, but also has efficiency implications. Yet, little is known about the fundamental drivers of teacher sorting across schools. Wage rigidity – a common feature of most public schooling systems around the world – has been highlighted as an important friction that would make teachers sort on non-pecuniary aspects of employment (Rosen, 1986), potentially resulting in low-quality teachers disproportionately working in disadvantaged schools (Clotfelter et al., 2005; Mansfield, 2015). While decentralizing wage-setting has been shown to yield limited equity and efficiency gains (Biasi et al., 2021), centralized, rule-based compensation reforms may be effective in addressing these challenges.

In this paper, we conduct a comprehensive analysis of the importance of teacher compensation design in shaping teacher sorting across schools and the distribution of student achievement. Our study takes place in Perú, a large developing country featuring stark geographic disparities in student achievement, a rigid teacher compensation schedule, and a centralized teacher recruitment drive. Using rich data on the universe of public-school teachers and students, we show that wage rigidity prompts teachers to sort on non-pecuniary aspects of employment, creating large urban-rural disparities in teacher effectiveness. We leverage a teacher compensation reform to show that raising wages in disadvantaged schools is effective at attracting qualified teachers and improving student learning. We then estimate a model of teacher sorting and student achievement featuring rich heterogeneity in

teachers' preferences and effectiveness. We allow for teacher-student match effects in our student achievement production function. We find that large equity and efficiency gains are attainable by either reallocating existing teachers to exploit teacher-student match effects or attracting applicants with higher average effectiveness into public teaching. Finally, we develop a counterfactual wage-setting procedure that seeks to achieve these gains by using information on teachers' preferences and effectiveness. Our results suggest that leveraging the extensive margin of recruitment is more cost-effective than inducing teacher reallocation to exploit match effects.

Our analysis draws on rich administrative panel data linking the universe of applicants and jobs posted in the nationwide centralized recruitment drive for public teachers in Perú with an array of additional data sources on public primary schools and students. All applicants take a standardized national competency test and choose their preferred position sequentially according to their test score rank. We first document that students in rural areas lack access to basic amenities and attend schools with notably inferior infrastructure. Teachers working in rural schools are significantly less qualified: they obtain 0.67 standard deviation (σ) lower scores in the national teacher competency test. Auxiliary survey evidence eliciting teachers' preferences over various non-wage job amenities suggests that the observed sorting is driven by preferences for urban amenities that are uncompensated by the rigid wage schedule. As a result, highly competent teachers disproportionately concentrate in urban areas. These inequities are associated with a staggering urban-rural difference of 0.55σ in students' standardized math scores.

We identify the causal effect of financial incentives on teacher sorting by evaluating a teacher compensation reform. This policy attributes wage bonuses to teachers working in rural schools that get increasingly large with remoteness through discrete jumps determined by arbitrary cutoffs on the school locality's population and its distance to the provincial capital. Using a regression discontinuity design around a 13% wage increase, we find that labor supply is highly elastic: higher wages raised new recruits' competency scores by 0.42σ . These gains passed through to students, raising school-level math and language scores by 0.27σ and 0.21σ , respectively. We provide evidence that these gains did not come at the expense of lowering the quality of teachers in lower-paying schools located near the threshold. Importantly, these effects are entirely driven by the selection of higher-quality new recruits rather than increased effort from incumbents.

We then build an empirical model of teacher sorting across schools and student achievement to characterize the potential equity and efficiency gains from alternative compensation policies. We allow teachers to have heterogeneous preferences over wages and non-wage job attributes. Specifically, non-pecuniary factors such as local amenities, geographical proximity, and teaching conditions induce both vertical and horizontal differentiation across jobs. In line with the institutional framework, wages are fixed, and we assume that the equilibrium teacher-school match is stable with respect to teachers’ preferences over schools and school priorities. Teacher sorting maps into student achievement through a potential outcomes framework where teacher effectiveness is allowed to be heterogeneous and interact with students’ prior achievement measures and gender. Teachers’ absolute and comparative advantages flexibly correlate with latent teacher attributes governing their willingness to pay for non-wage amenities and their valuation of the outside option. This potentially captures intrinsic motivation that cannot be explained by observable teacher characteristics and allows for selection on unobserved teaching quality in response to counterfactual teacher compensation policies.

Stability implies that each teacher is matched to their preferred school among their feasible choice set, i.e. the set of schools that would be willing to rematch with them. As school priorities are observed, we can construct the set of feasible schools for each teacher directly from the data. This unique feature allows us to express the equilibrium teacher-school allocation as the outcome of a discrete choice problem with personalized choice sets. We leverage this insight to identify and estimate teachers’ preferences from the observed teacher-school match. We relax the selection on observables assumption typically imposed in the teacher value-added literature by allowing teachers’ unobserved preference for their school assignment to directly influence student outcomes. Standard selection correction recovers unbiased estimates of the teacher value-added parameters.

The estimated model closely replicates the main features of the data, including the equilibrium sorting patterns and the causal effects on teacher quality and student achievement estimated at the RD threshold, as well as moments of the distribution of matched teacher and school characteristics. Our estimates show that teachers’ willingness to pay for non-wage amenities is typically large and features substantial heterogeneity. Consistent with the descriptive survey evidence, teachers have a high willingness to pay for proximity to home and better teaching conditions, driving the concentration of talent in urban areas. To fully

compensate for the urban-rural amenity differential, the most remote schools would require wage premiums two to four times larger than existing levels.

We report substantial heterogeneity in teachers’ absolute and comparative advantage. Moving one standard deviation up the distribution of average effectiveness implies an increase in student scores of 0.38σ in math and 0.33σ in Spanish. Importantly, 22-38% of the overall variance in teacher effectiveness can be explained by differences in their comparative advantage. The variance in teacher value-added is larger for students lagging behind, implying that rural schools would highly benefit from making teachers sort based on their comparative advantage. Teachers’ latent types driving their preferences over job postings correlate with their effectiveness. In particular, teachers who are more responsive to wage differences and those who have better outside options are less effective on average. Accounting for such selection is crucial to get accurate predictions of the distribution of student achievement under counterfactual compensation regimes.

Using our teacher value-added estimates, we show that, even in the presence of the wage bonus policy, teacher sorting across locations remains highly unequal, favoring urban areas. In particular, we document a 0.14σ urban-rural gap in teacher value-added for math, corresponding to one quarter of the overall gap in student achievement. Teachers do not sort based on their comparative advantage, implying large potential efficiency and equity gains from counterfactual teacher assignments. We quantify these gains by characterizing the allocation that maximizes total student achievement. Efficiently reallocating the pool of assigned teachers would increase student test scores in rural areas by 0.17σ while moderately decreasing student test scores in urban areas by 0.02σ . In contrast, inducing applicants with higher average effectiveness to sort into public teaching instead of choosing the outside option could generate aggregate gains of 0.11σ .

Finally, we provide a counterfactual wage-setting protocol that seeks to achieve these equity and efficiency gains by leveraging information on teachers’ preferences and effectiveness. We construct an iterative wage-adjustment procedure that mirrors the wage-proposal dynamics of generalized Deferred Acceptance in matching-with-contracts environments ([Hatfield and Milgrom, 2005](#)). The procedure iteratively increases compensation at rural schools that fail to attract a teacher with sufficiently high value-added and recomputes the equilibrium assignment until the policy objective is attained. Applying this procedure, we show that the status quo policy is poorly targeted and that large equity and efficiency gains are

attainable at no additional cost. Importantly, these gains mainly come from attracting applicants with higher average effectiveness into public teaching rather than reallocating existing teachers to exploit teacher-student match effects. This suggests that, despite the large potential gains from teacher reallocation, leveraging the extensive margin of recruitment is more cost-effective.

Our findings speak to several strands of literature. A large body of work studies teachers' contribution to student achievement (Rivkin et al., 2005; Chetty et al., 2014b; Araujo et al., 2016; Bau and Das, 2020). We first highlight the benefits of leveraging data on teachers' choices over schools to identify and estimate teacher value-added. As teachers' preferences over job attributes correlate with their effectiveness, this additional information is essential to capture selection on unobserved quality resulting from changes in compensation (Rothstein, 2015; Brown and Andrabi, 2020). Following Abdulkadiroğlu et al. (2020), we also show how to leverage this information to correct for selection on teachers' unobserved preference for their school assignment. Finally, we add to the literature studying teacher-student match effects (Ahn et al., 2025) by showing that efficiently reallocating teachers would largely reduce urban-rural disparities in student achievement.

This paper also contributes to a recent literature studying the link between teacher sorting and student achievement through equilibrium models of the labor market for teachers (Boyd et al., 2013; Tincani, 2021; Biasi et al., 2021; Bates et al., 2025). In contrast to the literature, our approach utilizes panel data on a nationwide centralized allocation mechanism for public-sector teachers, providing detailed information on each applicant's choices and choice sets over several years. This dataset, combined with quasi-experimental variation in wages, allows us to identify and estimate, under minimal assumptions, a rich empirical model of teacher sorting and student achievement accounting for flexible heterogeneity in teachers' preferences and effectiveness.

Finally, our work is broadly related to a growing literature studying personnel and organizational policies in the public sector (Finan et al., 2017). While there is a large body of work studying the effectiveness of pay-for-performance schemes for teachers (Muralidharan and Sundararaman, 2011; Barrera-Osorio and Raju, 2017; Biasi, 2021; Leaver et al., 2021; Gilligan et al., 2022; Brown and Andrabi, 2020), there is relatively little evidence on the effects of unconditional pay increases on teacher sorting and student outcomes (Clotfelter et al., 2008; de Ree et al., 2018). In line with this literature, we find that such interventions

do not prompt increased effort from incumbent teachers. However, we document that they can largely increase student achievement by attracting higher-quality teachers. We borrow insights from the empirical market design literature ([Agarwal, 2017](#); [Agarwal and Budish, 2021](#)) to develop a wage-setting procedure aimed at attaining equity and efficiency gains by leveraging information on teachers’ preferences and effectiveness. Our results show that targeted compensation can generate substantial gains in student achievement at no additional cost by attracting applicants with higher average effectiveness into public teaching.

2 Context, Data, and Descriptive Evidence

2.1 Institutional Background

Our analysis focuses on public primary education, which offers the broadest coverage of any schooling level in Perú. Secondary schools are notably less prevalent in rural regions. Public schools constitute 75% of nationwide primary school enrollment in Perú. In rural areas, more than 26,000 public primary schools enroll 99% of school-aged children.

Public school teachers in Perú are hired under two distinct types of contracts. Permanent teachers (*docentes nombrados*) are civil servants with secure employment conditions. Contract teachers (*docentes contratados*) are hired on a fixed one-year contract by a specific school, renewable for an additional year upon approval from the school’s principal.

Primary school teachers’ earnings in Perú rank second to last among liberal professions ([INEI, 2016](#)). In 2016, all contract teachers were receiving a fixed base monthly wage of S/ 1,396 (US\$ 402) while permanent teachers were receiving S/ 1,550 (US\$ 447), irrespective of where they worked.¹ Public-sector teachers also receive wage bonuses linked to specific school appointments. Figure 1 illustrates the various wage bonuses that were in place during our analysis period, which range between 4% (for bilingual schools) and 36% (for extremely rural locations, as detailed in Section 3.1) of the monthly base wage. These bonuses are additive, such that teachers working in schools meeting multiple criteria (e.g., being both multi-grade and rural) accumulate bonuses.

Historically, the recruitment of public teachers followed a decentralized approach, granting regional and local officials substantial discretion in hiring and resource allocation. To

¹This figure increased to S/ 2,000 in 2018. The average wage for a primary-school teacher in the private sector was S/ 950 (US\$ 274) per month ([MINEDU, 2014](#)).

Figure 1: Wage Bonuses for Teachers in Disadvantaged Schools



NOTES. This figure shows the monetary amount in Peruvian Soles (S/) for the different wage bonuses implemented by the Government during the period 2015-2018. The VRAEM indicator groups schools that are in the *Valle de los Rios Apurimac, Ene y Mantaro*, an extremely poor region that presents security concerns due to the activity of drug cartels. Border categorizes schools that are adjacent to the country's borders. See Section 3.1 for the definition of Extremely Rural, Rural, and Moderately Rural.

enhance transparency and fairness, the Ministry of Education (MoE) introduced a nationwide recruitment process in 2015, which centralized all job postings and teacher applications in a unified platform. This process has taken place every other year since then. All applicants for positions in public schools are required to have a teaching certification (i.e. a teacher degree) and to undergo a national competency test.² Permanent positions are solely accessible to teachers scoring above selective thresholds and are assigned through a two-step matching process in which schools have a fair amount of discretionary power over hires.³ In our data, only about 11% of the applicants are eligible for a permanent teaching position. Temporary positions are allocated through a serial dictatorship algorithm where applicants are ranked according to their competency score. Applicants first select a school district and then choose from the available vacancies within that district by order of their rank.

We restrict the analysis to the recruitment of contract teachers for several reasons. First, the majority of vacancies open in rural schools are filled by teachers on a temporary contract (85%), thus making the allocation of short-term teaching positions most relevant for student

²The test comprises three modules, each contributing different weights to the total score: logical reasoning (25 percent), reading comprehension (25 percent), and curricular and pedagogical knowledge (50 percent).

³Eligible applicants can indicate their preferred school district and submit a rank order list of schools within that district. In a second evaluation round, teachers are given another score based on their CV, an in-person interview, and a classroom demonstration.

outcomes in remote locations. Second, the recruitment of contract teachers is based on a “strict-priority” mechanism – schools rank candidates by their competency score. This institutional feature enables us to observe teachers’ choice sets directly, which is essential to identify teachers’ preferences over job postings.

2.2 Data Description and Sample Selection

Our empirical analysis draws upon comprehensive administrative records on students, teachers and schools within the Peruvian public education system. Through the centralized matching platform, we have access to information on the universe of vacancies and applicants for teaching positions at public schools in 2016 and 2018. These data include the wage associated to each job posting, applicants’ scores in the national competency test, and basic demographics such as education, age, gender, and native language.

We complement this information with several auxiliary data sources providing a wide range of additional school and applicant characteristics. The school census collects yearly information on the number of students and teachers in each school, school infrastructure (libraries, computers, sports facilities, etc.), and access to basic local amenities like electricity, water, internet, banking services, and public libraries. The teachers’ employment records (NEXUS, for its acronym in Spanish) track all employed public-sector teachers from 2012 to 2020 and contain information on their workplace and contract type. Finally, we obtain information on the residential location of applicants from a government-collected dataset at the household level, which is primarily utilized for targeting social programs (SISFOH, for its acronym in Spanish).

Student academic performance is obtained from a national, standardized test evaluating proficiency in math and Spanish language (ECE, for its acronym in Spanish). This assessment has near-universal coverage of the student population in the country and contains individual test scores for second- and fourth-grade students. We can link students to their respective classroom and teacher identifiers through an administrative teacher-classroom dataset held by the Ministry of Education (SIAGIE, for its acronym in Spanish).

We construct two samples for our analysis. The *choice sample*, used to analyze teachers’ choices over job postings, is composed of applicants and positions that participated in the 2016 and 2018 recruitment drives for contract teachers. Considering only schools and applicants with non-missing covariates, our choice sample comprises 22,743 vacancies and

50,319 applicant-year observations.⁴ Importantly, 75% of applicants participating in the 2016 recruitment drive reapply in 2018, which allows us to construct a panel tracking teachers’ re-matching decisions. We label our second analysis sample as the *outcome* sample, which we use to measure teacher effectiveness and to study the effects of recruitment on student achievement. This is a subset of the choice sample composed of applicants assigned to fourth-grade classrooms for which we observe student performance at the national test. Our outcome sample comprises 4,834 teachers and 96,508 students. Table A.1 provides descriptive statistics for these two samples.


2.3 Wage Rigidity, Sorting, and Inequalities

Educational outcomes in Perú exhibit severe spatial inequality. Despite policy efforts to address historical underinvestment in disadvantaged areas (Bertoni et al., 2020), the urban-rural divide remains acute. This disparity is particularly evident in the stark gaps in both schooling inputs and student performance. While urban amenities are nearly universal, rural localities frequently lack basic facilities, such as internet access, medical clinics, or access to drinking water. School-specific infrastructure is similarly stratified; for instance, only 8% and 37% of rural schools have science-lab and computer facilities, while the corresponding shares in urban schools are 42% and 83%, respectively. These disparities, along with differences in other school characteristics, are reported in Table A.2.

Figure 2 visually portrays the striking disparities in teaching competency and student achievement between primary schools in rural and urban areas. As shown in Panel A, teachers working in rural schools score, on average, 0.67σ lower on the national competency test than teachers working in urban schools. Teachers in the most remote areas, categorized as *Extremely Rural* (see Section 3.1 for a formal definition), score 0.80σ lower than in urban schools. These spatial differences strongly correlate with inequalities in student achievement. Panel B displays the distribution of students’ academic performance at the national standardized evaluation. We observe an urban-rural gap in student achievement of 0.55σ in mathematics and 0.63σ in Spanish.

⁴When removing a matched applicant because of missing information, we remove the corresponding position. Similarly, when removing a matched position, we remove the corresponding applicant. Under this sampling procedure, the likelihood of removing matched vacancies and applicants is twice as large as the likelihood of removing unmatched vacancies and applicants (Menzel, 2015). To correct for this and ensure that the overall matching rate remains constant, we randomly remove unmatched applicants and vacancies accordingly.

Figure 2: Teacher Competency and Student Achievement by Remoteness



a) Teacher Competency by Urban/Rural

b) Student Achievement by Urban/Rural

NOTES: This figure presents summary statistics for teachers and students in schools with different levels of rurality. Areas are grouped as follows: (a) urban areas, where the locality population exceeds 2,000 inhabitants; (b) extremely rural areas, where the locality population is below 500 inhabitants and the travel time to the closest provincial capital exceeds 120 minutes; and (c) other rural areas, where the locality population is either above 500 inhabitants or travel time is below 120 minutes (see Section 3.1 for details). Panel A uses our choice sample and reports the average teacher competency score in filled vacancies by rurality. Panel B uses our outcome sample and reports the average student scores in the fourth-grade Spanish and math modules of the national standardized assessment by rurality. In each panel, diamonds indicate the average for urban and rural areas, where the rural average aggregates the two rural categories.

Several factors might explain the observed teacher sorting patterns. Teachers working in rural areas contend with several challenges, including the scarcity of basic school resources, inadequate services, limited access to public goods, and, for many, geographical distance from home. We directly elicit teachers' preferences using an online survey among applicants for permanent positions during the 2016 centralized job application process, in which we obtained a response rate of just under 20% (5,550 teachers).⁵ The survey results indicate that non-monetary factors significantly influence teachers' choices over job postings. As shown in Table 1, 44% of teachers mention 'proximity to home' as a crucial factor guiding their preference ranking. Likewise, attributes such as prestige, safety, and cultural considerations are also frequently cited as relevant when ranking potential teaching positions.

Since rigid wages fail to compensate for rural disamenities, teachers with higher priority systematically sort into urban schools. As a result, over half of urban postings (compared to a quarter in rural areas) are occupied by teachers ranked in the top 20% of the applicant pool

⁵As shown in Table A.3, observable teacher characteristics of survey respondents align closely with the population of applicants for permanent positions. Test scores of survey respondents are significantly higher than those within our choice sample (see Table A.1) since the survey was sent to teachers who were eligible to apply for permanent teaching positions, i.e. who scored above 120 points in the test.

Table 1: Applicant Survey (Choice Attributes)

	Rank			In Top 3
	1 st	2 nd	3 rd	
Close to House	44.17	11.66	8.00	63.83
Safe	10.66	24.19	19.25	54.1
Well Connected	9.69	20.62	20.20	50.51
Prestige	17.92	14.12	12.29	44.33
Cultural Reasons	10.61	9.67	12.31	32.59
Good Infrastructure	2.02	8.40	12.86	23.28
Good Students	1.24	4.52	6.08	11.84
Possibility other Jobs	1.93	3.72	4.90	10.55
Career	1.76	3.10	4.09	8.95

NOTES. This table summarizes the answers of 5,550 survey respondents to the question “What are the most important characteristics for your ranked choices?” Survey participants are applicants for permanent positions in the 2016 application process. The first three columns show the share of respondents that ranked the corresponding answer first, second, or third. The last column shows the share of respondents who listed the corresponding choice in their top 3 reasons.

within their school districts. This suggests that wage rigidity is an important contributor to unequal teacher sorting and the associated spatial disparities in student outcomes. In Section 3, we study the scope for teacher compensation reform to address these disparities by leveraging a policy that provided wage bonuses to teachers working in rural schools.


3 Teacher Compensation Reform

3.1 The Rural Wage Bonus Policy

Wage bonuses for contract teachers working in rural schools were introduced in August 2015, i.e., during the school year before the first wave of the centralized teacher recruitment drive. The reform was only announced briefly before being actually implemented. Both incumbent and newly assigned teachers were eligible for these bonuses. The policy established three categories of “rurality”, which are defined based on two measures: (i) the school locality’s population and (ii) travel time between the school and the provincial capital. The population of the locality is measured by the latest available census (2007), while travel time to the provincial capital is computed based on the school’s GPS coordinates (taken on-site by government inspectors), the type of roads available, and the most frequent modes of transport.

Figure 3 shows how wage bonuses are distributed across schools based on the two measures

Figure 3: The Distribution of Rural Schools and the Wage Bonuses



NOTES. This figure shows the spatial distribution of rural primary schools along the two dimensions that determine the assignment of the rural wage bonus: travel time to the closest provincial capital (y-axis) and locality population (x-axis). The solid black lines indicate the cutoff(s) defining the rurality categories. Each category's name, together with the corresponding wage bonus (in Peruvian soles), is reported in the labeled boxes. The sample includes all schools in the choice sample.

defining the rural categories. *Extremely Rural* schools are in localities with less than 500 inhabitants and situated more than 120 minutes away from the provincial capital. Teachers in these schools receive a bonus of S/ 500 (US\$ 144), representing between 25% and 36% of their base wage (depending on the year of the assignment).⁶ *Rural* schools are either in localities with less than 500 inhabitants and situated between 30 and 120 minutes away from the provincial capital, or in localities with 500 to 2,000 inhabitants that are farther than 120 minutes from the provincial capital. The bonus received by teachers in these schools is S/ 100 (US\$ 29). Finally, *Moderately Rural* schools are either in localities with 500 to 2,000 inhabitants that are within 120 minutes of the provincial capital or in localities with less than 500 inhabitants that are within 30 minutes of the provincial capital. In these schools, teachers receive a bonus of S/ 70 (US\$ 20). All other schools are classified as urban and are therefore not entitled to the rural wage bonus.

There is a large mass of schools around both the travel time (30 minutes and 120 minutes from the provincial capital) and the population cutoffs (500 inhabitants). As localities become

⁶The base monthly wage of contract teachers increased from S/ 1,396 in 2016 to S/ 2,000 (US\$ 576) by the end of 2017.

more remote, they are more likely to have few inhabitants and predominantly fall into the *Extremely Rural* category. Likewise, as localities become more populated, they are less likely to be remote and fall into the *Moderately Rural* category.

3.2 Regression Discontinuity Design

To study the effects of increasing compensation in rural schools on teacher sorting and student achievement, we exploit the sharp thresholds that determine the allocation of the rural wage bonuses in a regression discontinuity (RD) design. The validity of this research design relies on two assumptions: (i) continuity of potential outcomes around the cutoffs, and (ii) independence between the potential outcomes of each unit and the treatment status of other units in a neighborhood of the cutoffs (SUTVA).

Continuity may be violated if the introduction of the rural wage bonus generated incentives for school administrators to manipulate the information used to determine eligibility for the bonus. The population cutoff of 500 inhabitants is based on census data collected before the policy was announced, and as such, is impossible to manipulate. The travel time cutoffs at 30 minutes and 120 minutes are based on GPS measures gathered periodically by government inspectors to account for possible changes in the transportation network and could be subject to manipulation. Figures [A.2](#) and [A.3](#) show a large and significant jump in the density of schools located just above the travel time threshold at 120 minutes in 2018 (but not for 2016) in both of our analysis samples. Instead, there are no significant jumps in the density of schools at the population threshold for either year for our two analysis samples.

SUTVA may be violated if the policy triggered spillovers through teacher sorting *around* the population cutoff – e.g. if teachers who chose a position in a high-bonus school just below the threshold would have otherwise chosen a position just above the threshold in the absence of the wage bonus policy. While we cannot entirely rule out this possibility, in [Section 3.3](#) we report evidence that suggests that spatial spillovers of this sort are not a first-order concern in our setting.

Given the possible manipulation of the travel time threshold, we only exploit the variation in wages generated by the population threshold in this part of the analysis. The validity of the RD design around the population threshold is corroborated by covariate smoothness tests

on both samples (see Table A.4).⁷ Teachers recruited in localities with slightly less than 500 inhabitants earn, on average, about S/ 244 (US\$ 72) more than teachers hired in localities that are just above the cutoff (corresponding to an increase of about 13%).⁸

To study the effects of the wage bonus policy on recruiting, we consider the following RD specification:

$$y_{ijt} = \gamma_0 + \gamma_1 \mathbf{1}(pop_j < pop_c) + g(pop_j, pop_c) + \delta_t + \varepsilon_{ijt}, \quad (1)$$

where y_{ijt} is the recruiting outcome for vacancy i (e.g. the competency score of the hired teacher) in school j at time t , $g(\cdot)$ is a flexible polynomial in the population of the locality of the school on both sides of the population cutoff, δ_t denotes time indicators for the specific year of the recruitment drive, and ε_{ijt} is an error term clustered at the school-year level.

We further use Equation (1) to study the effects of the wage bonuses on student outcomes, denoted as $y_{\ell jt}$, e.g. student ℓ 's standardized test score in school j at time t . In both analyses, the parameter of interest is γ_1 , which represents the average outcome difference between vacancies or students in localities that are just below or above the population cutoff. We estimate γ_1 non-parametrically using the robust estimator proposed by Calonico et al. (2014) through bias-corrected local linear regressions that are defined within the mean squared error optimal bandwidths.

We exclude all urban schools and rural schools in localities within 30 minutes of the provincial capital, since crossing the population cutoff does not lead to an increase in the rural wage bonus in these schools (see Figure 3). In what follows, we refer to the observations in rural areas that are above the population cutoff of 500 inhabitants as ‘Low-Bonus’ and to those below the cutoff as ‘High-Bonus’.


3.3 Teacher Sorting

We start by investigating how teachers’ school choices responded to the wage increase at the threshold using our choice sample. The graphical evidence displayed in Panel A in Figure 4,

⁷Table A.4 shows that the indicator variable for the travel time threshold (travel time >120) does not significantly jump at the population threshold. This addresses concerns that schools might have manipulated the 120-minute threshold to a greater extent below the 500-inhabitant threshold (where the benefits from manipulating are larger), which would threaten the validity of our research design.

⁸Note that the wage increase at the threshold would be larger if we restrict our estimation sample to schools above the travel time cutoff of 120 minutes (see Table A.5). However, this alternative approach would imply conditioning on a partially manipulated variable (travel time) and decrease the sample size.

Figure 4: Teacher Choices and Sorting



NOTES. This figure shows how the probability that a vacancy is filled (Panel A) and the competency score of the assigned teacher (Panel B) vary with the population of the school's locality on either side of the 500-inhabitant cutoff. Each marker indicates the mean of the outcome variable within each bin, defined following the IMSE-optimal evenly spaced method proposed by [Calonico et al. \(2015\)](#). Solid lines represent predictions from linear regressions estimated separately on either side of the cutoff, assuming a triangular kernel. Dashed lines denote 95% asymptotic confidence intervals. The sample used is the choice sample. Additionally, the sample excludes urban schools (i.e., those located in areas with more than 2,000 inhabitants) and schools situated within 30 minutes of the nearest provincial capital. In Panel B, the sample is further restricted to vacancies filled through the centralized assignment process.

along with the corresponding RD estimates reported in Column (1) of Table 2, documents that the wage increase at the threshold had a small and statistically insignificant effect on the probability that a vacancy is filled in the national recruitment drive.

We next restrict our attention to filled vacancies on either side of the threshold and consider as outcome variables a measure of teachers' preference intensity and the competency score of the assigned teachers.⁹ The preference index takes the value of zero if the position is filled last and the value of one if the position is filled first in a school district. We find that high-paying vacancies were filled at a significantly higher priority order within the centralized assignment process than low-paying vacancies. The average vacancy in a low-bonus location is filled by a teacher ranked in the 66th percentile of the score distribution of applicants in that school district. Column (2) of Table 2 shows that high-bonus schools fill the position with an applicant ranked in the 77th percentile on average.

Given that priority for choosing vacancies in the recruitment drive for short-term positions is determined by teachers' competency scores only, we expect that the increase in preferences

⁹These outcomes are observed for filled vacancies only, raising concerns about selection. We thus report in Table A.6 bounds around our RD estimates for both of these outcomes using the approach outlined in [Gerard et al. \(2020\)](#). The bounds are tight, suggesting that the potential selection caused by restricting the analysis to filled vacancies is inconsequential for the results of our RD analysis.

Table 2: Teacher Choices and Sorting

	(1) Vacancy filled	(2) Preferences	(3) Teacher Score (Std.)
High Bonus	0.063 (0.048)	0.108 (0.031)	0.418 (0.106)
Bandwidth	190.896	162.270	193.380
Schools	1478	1157	1388
Observations	3524	2575	3068

NOTES. This table reports the effect of crossing the population threshold on several outcomes. In Column (1), the outcome variable is an indicator for whether the vacancy is filled in the national recruitment drive. Column (2) reports the effect on the rank at which a vacancy is chosen within a school district, normalized to range from zero to one. Column (3) reports the results on the competency score obtained by teachers in the centralized test. In all regressions, we use our choice sample and exclude urban schools (i.e., those located in areas with more than 2,000 inhabitants) and schools situated within 30 minutes of the nearest provincial capital. In Columns (2) and (3), the sample is further restricted to vacancies filled in the national recruitment drive. Cells report bias-corrected regression-discontinuity estimates obtained using the robust estimator proposed by [Calonico et al. \(2014\)](#). Regressions are estimated within a mean-squared-error-optimal bandwidth, reported at the bottom of the table. Standard errors are clustered at the school $\times$ year level.


for high-bonus schools documented in Column (2) translates into an increase in the quality of recruited teachers. Both the graphical evidence in Panel B of Figure 4 and the RD estimates in Column (3) of Table 2 document that teachers who sort into high-bonus schools have higher competency scores than those who choose a position in low-bonus schools. The effect is sizable, at 0.42σ of the overall distribution of the competency score. As a benchmark, this corresponds to approximately twice the gap in teacher competency between *Extremely Rural* schools and other rural schools and two-thirds of the overall urban-rural gap (see Figure 2).

In sum, increasing wages in disadvantaged locations effectively steered teachers' labor supply toward the targeted job postings. The observed change in teachers' behavior does not seem to have significantly affected the probability of creating new matches, but instead led to an inflow of more competent teachers in high-bonus schools.¹⁰

Additional evidence suggests a limited role of spillovers around the 500 population threshold. First, it is important to note that localities around the population threshold are not necessarily geographically close to each other. In fact, for any given school below the cutoff, the median (geodesic) distance to its first, second, and third closest school above the cutoff is approximately 10km, 20km, and 30km, respectively. Second, we use information on teachers' previous assignments to quantify potential reallocation between low-bonus and high-bonus schools around the threshold induced by the policy. Figure A.4 shows no evidence that teachers working in high-bonus schools are more likely to have previously worked

¹⁰This result is consistent with recent evidence reported in [Agarwal \(2017\)](#), which documents that the primary effect of financial incentives was to increase the quality, not numbers, of medical residents in rural America.

Figure 5: Wage Bonus and Student Achievement



NOTES. This figure shows how student achievement varies with the population of the school's locality on either side of the 500-inhabitant cutoff. Panel A reports standardized test scores in Spanish, while Panel B reports standardized test scores in mathematics. Each marker indicates the mean of the outcome variable within each bin, defined following the IMSE-optimal evenly spaced method proposed by [Calonico et al. \(2015\)](#). Solid lines represent predictions from linear regressions estimated separately on either side of the cutoff, assuming a triangular kernel. Dashed lines denote 95% asymptotic confidence intervals. The sample used is the outcome sample, focusing on students taught by newly recruited teachers (corresponding to Column (1) of Table 3). Additionally, the sample excludes urban schools (i.e., those located in areas with more than 2,000 inhabitants) and schools situated within 30 minutes of the nearest provincial capital.

in low-bonus schools close to the threshold. The estimated effect for schools just above the threshold is negative, suggesting that high-bonus schools are less likely to draw teachers from schools located just above the population cutoff relative to the rest of the locality-size distribution. Finally, Table A.7 reports the results of a difference-in-discontinuity analysis, which suggests that the increase in teacher competency in high-bonus schools did not come at the expense of decreasing the competency scores of teachers assigned in low-bonus schools near the threshold.

3.4 Student Achievement

Offering higher wages for positions in rural locations could potentially improve student outcomes through two mechanisms: (i) the recruitment of higher-quality teachers, and (ii) an increased effort from incumbent teachers. To parse out the effort margin from the recruitment margin, we leverage matched teacher-classroom data (i.e. our outcome sample) and run the analysis separately for classrooms taught by newly recruited teachers and for those taught by incumbent teachers. We use fourth-grade student test scores in math and Spanish in 2016 and 2018 as measures of academic achievement.

Table 3: Wage Bonus and Student Achievement

<i>Panel A: Dependent Variable is Spanish Test (z-score)</i>			
	(1)	(2)	(3)
	New teacher	No vacancy	All
High Bonus	0.340 (0.174)	-0.080 (0.124)	0.206 (0.084)
Bandwidth	92.657	143.045	114.510
Schools	221	436	931
Observations	3830	6578	18828
<i>Panel B: Dependent Variable is Math Test (z-score)</i>			
	(1)	(2)	(3)
	New teacher	No vacancy	All
High Bonus	0.543 (0.221)	-0.040 (0.143)	0.265 (0.099)
Bandwidth	79.168	144.756	110.039
Schools	192	442	889
Observations	3408	6699	18083

NOTES. This table reports the effect of crossing the population threshold on student achievement in Math and Spanish. The outcome variables are standardized test scores in Spanish (Panel A) and Math (Panel B) for fourth-grade students. Column (1) considers the sample of fourth-grade students in our outcome sample (i.e. taught by newly recruited contract teachers). Column (2) considers schools with no open vacancies for contract teaching positions. Column (3) considers the overall sample of students, irrespective of whether they are taught by a newly assigned or an incumbent teacher. In all columns, the sample excludes urban schools (i.e., schools located in areas with more than 2,000 inhabitants) and schools situated within 30 minutes of the nearest provincial capital. Each cell reports bias-corrected regression-discontinuity estimates obtained using the robust estimator proposed by [Calonico et al. \(2014\)](#). Regressions are estimated within a mean-squared-error-optimal bandwidth, reported at the bottom of the table. Standard errors are clustered at the classroom $\times$ year level.

Figure 5 plots the threshold crossing effects on students’ academic achievement in classrooms taught by newly recruited teachers. Students exposed to higher-quality teachers, attracted by the wage increase at the threshold, perform significantly better on standardized tests. Column (1) in Table 3 shows sizable effects ranging from 0.34σ in Spanish to 0.54σ in math. We find no evidence of an increased effort from incumbent teachers in response to the wage increase at the threshold. Column (2) of Table 3 shows very small effect sizes of the wage bonus when restricting the sample to schools with no open vacancies. This result mirrors recent findings that establish little or no effects of unconditional wage increases in contexts where most teachers are public servants with permanent contracts ([de Ree et al., 2018](#)). Finally, Column (3) displays the aggregate effect on the overall sample, including students taught by new recruits and incumbents. Students in high-bonus schools perform overall better in Spanish and math, with effect sizes of 0.20 - 0.26σ .

This evidence shows that the teacher compensation reform attracted more effective teachers at the targeted locations, which in turn improved students’ learning outcomes. We next build a model of teacher sorting and student achievement to quantify the equity and efficiency

gains attainable by counterfactual teacher compensation design.

4 An Empirical Model of Teacher Sorting and Student Achievement

4.1 Wages, Preferences, and Equilibrium

We consider a population of N teachers indexed by i , J schools indexed by j , and L students indexed by ℓ . Time is indexed by t . We denote the realized school assignment of teacher i in year t as $D_{it} \in \{0, 1, \dots, J, J + 1\}$.

Wages are set by the government through a known deterministic rule. They are posted ex-ante for each available vacancy, observed by applicants before making their choices, and they cannot be renegotiated. Teachers receive the same fixed baseline wage, irrespective of the school they work in, and a set of wage bonuses, which vary with pre-determined locality and school characteristics (see Figure 1). We denote the monthly wage (including any applicable bonuses) in school j in year t as w_{jt} .

We define the indirect utility that teacher i gets from being matched with school j in year t as:

$$\begin{aligned} U_{ijt} &= u_{ijt} + \eta_{ijt} \\ &= \alpha_i w_{jt} + u(x_i, z_{jt}, d_{ijt}) + \kappa_{c(j)} + \eta_{ijt}, \end{aligned} \tag{2}$$

where $u(x_i, z_{jt}, d_{ijt})$ is a flexible function of observed school characteristics (z_{jt}), teacher characteristics (x_i) and measures of geographical proximity or familiarity between school j and teacher i (d_{ijt}) giving rise to vertical and horizontal differentiation between jobs offering the same wage, $\kappa_{c(j)}$ are province fixed effects capturing vertical differences in unobserved non-pecuniary factors across locations, and η_{ijt} is an unobserved taste shock introducing further horizontal differentiation in teachers' preferences. The parameter α_i controls teacher i 's taste for wages relative to other non-pecuniary factors.

We specify the function u as follows:

$$u(x_i, z_{jt}, d_{ijt}) = x_i' \Theta_1 z_{jt} + x_i' \Theta_2 d_{ijt}$$

where the vector z_{jt} contains an amenity/infrastructure index,¹¹ indicators for whether school j belongs to specific categories that are eligible for other wage bonuses such as multi-grade, single-teacher, bilingual, and/or specific geographic areas (see Figure 1), polynomials of the locality’s population and travel time (in hours) to the provincial capital,¹² and school averages of student characteristics including their 2nd grade test scores in math and Spanish, gender, age, and ethnolinguistic background. The vector d_{ijt} contains linear splines of the geodesic distance between school j and teacher i ’s home location, and an indicator for whether teacher i was working in school j in the previous year. The vector x_i contains an intercept, gender, competency score, indicators for experience in the public and private sector, and ethnolinguistic background.

Around two-thirds of the applicants remain unassigned at the end of the centralized assignment procedure. Among the unassigned applicants, half find alternative teaching positions within the public sector in later decentralized rounds. We denote this alternative as $j = J + 1$ and normalize its utility as follows:

$$u_{iJ+1t} = \gamma_1 \mathbb{1}\{t = 2018\} + x_i' \gamma_2, \quad (3)$$

where $x_i' \gamma_2$ captures heterogeneity across teachers in the value of participating in the secondary decentralized market.

The remaining unassigned applicants seek opportunities outside the public sector (e.g., teaching in a private school, or any other occupation, or staying outside of the labor force). We denote this alternative as $j = 0$ and normalize its utility as follows:

$$u_{i0t} = \beta_1 \mathbb{1}\{t = 2018\} + \beta_i, \quad (4)$$

where β_i is a teacher-specific random coefficient capturing observed and unobserved heterogeneity in the pecuniary and non-pecuniary benefits of choosing to leave the public sector. We assume that options $j \in \{J + 1, 0\}$ are both in the choice sets of all teachers.

¹¹The amenity index is constructed by aggregating different indicators measuring the quality of infrastructure in the locality (see Panel B of Table A.2) as well as an asset-based measure of poverty at the individual level that we aggregate at the locality level (which enters with a negative sign in the overall index) by PCA.

¹²Conditional on z_{jt} , the residual variation in wages w_{jt} stems from the discrete changes at the thresholds induced by the rural wage policy (see Section 3.1).

We define $\theta_i = (\log \alpha_i, \beta_i)$ and assume that

$$\theta_i | x_i \sim \mathcal{N}(\mu'_\theta x_i, \Sigma_\theta).$$

These correlated random coefficients capture time-invariant idiosyncratic factors that affect how teachers substitute non-pecuniary benefits against wages, as well as how they value the outside option. Allowing for these coefficients to be correlated is crucial to correctly pin down substitution patterns from the outside option in response to changes in wages across job postings. Observing two choices per teacher and having variation in choice sets within and across teachers helps in the identification of the distribution of random coefficients (Berry et al., 2004). For example, teachers choosing the outside option (4) in one year and a public school in another year provide essential variation to identify the off-diagonal elements of Σ_θ .

As described in Section 2.1, applicants are ranked based on their competency score within school districts and choose, by order of priority, their preferred school among those that still have open vacancies. This precludes the existence of blocking pairs, i.e. teachers would not be accepted by a school they strictly prefer to the one they choose (Roth and Sotomayor, 1992). As teachers choose their district before knowing which options are available to them, they might realize ex-post that they would have preferred a feasible school in another district. However, only 0.60% of unassigned applicants chose a position in another district in later decentralized rounds. We take this as evidence that justified envy is minimal in this setting and thus assume that the realized teacher-school assignment $D_{1t}, \dots, D_{Nt}$ is stable both within and across school districts. As a result, we can characterize the equilibrium teacher-school match as follows:

$$D_{it} = \arg \max_{j \in \Omega(s_{it}) \cup \{0, J+1\}} U_{ijt}, \quad i = 1, \dots, N, \quad (5)$$

where $\Omega(s_{it})$ is teacher i 's feasible choice set: the set of schools that still have remaining vacancies after all applicants with a score larger than s_{it} made their choice. As $\Omega(s_{it})$ is directly observed in the data, the matching equilibrium can be rewritten as a multinomial logit model with personalized choice sets (Fack et al., 2019).

We estimate the preference parameters $(\Theta_1, \Theta_2, \gamma_1, \gamma_2, \beta_1, \kappa)$ as well as the distribution of the random coefficients $(\mu_\theta, \Sigma_\theta)$ by maximizing the log-likelihood of observing the matching

history $(D_{it})_{t=1}^T$ for $i = 1, \dots, N$

$$\sum_{i=1}^N \log \int \prod_{t=1}^T \frac{\exp\{u_{iD_{it}t}\}}{\exp\{u_{i0t}\} + \exp\{u_{iJ+1t}\} + \sum_{k \in \Omega(s_{it})} \exp\{u_{ikt}\}} \phi(\theta_i | \mu'_\theta x_i, \Sigma_\theta) d\theta_i, \quad (6)$$

where we approximate the integral using 50 draws from a Halton sequence (Judd, 1998).

4.2 Student Achievement

We specify an education production function that maps potential teacher-student matches to student outcomes. In particular, we define $Y_{\ell i}$ as the potential outcome of student ℓ when taught by teacher $i \in \{1, \dots, N\}$ such that

$$Y_{\ell i} = \delta_{0i} + X'_{\ell 1} \delta_{1i} + X'_{\ell 2} \delta_2 + \epsilon_{\ell i}, \quad (7)$$

where $\mathbb{E}[\epsilon_{\ell i}] = 0$, $\mathbb{E}[X_{\ell 1} \epsilon_{\ell i}] = 0$ and $\mathbb{E}[X_{\ell 2} \epsilon_{\ell i}] = 0$ by construction. The vector $X_{\ell 1}$ includes student ℓ 's lagged test scores (second grade) in the corresponding subject and gender, and $X_{\ell 2}$ includes the variables in $X_{\ell 1}$ averaged at the classroom level, as well as squared lagged test scores, ethnicity, and age both at the student level and averaged at the classroom level. The vector $X_{\ell 2}$ also includes the school characteristics $z_{j(\ell)}$ entering teachers' preferences (2), where $j(\ell)$ denotes the identity of the school where student ℓ is enrolled.¹³ We normalize $X_{\ell 1}$ to be mean zero such that δ_{0i} can be interpreted as the average treatment effect of teacher i , while δ_{1i} captures the returns to $X_{\ell 1}$ when being exposed to teacher i . This allows teachers to have comparative advantages in their ability to teach students with different demographics or backgrounds.¹⁴ Accounting for such teacher-student match effects is particularly important for our counterfactual analysis, as it enables us to quantify the potential efficiency gains from alternative teacher assignments.

We denote the realized assignment of student ℓ to teachers as $S_\ell \in \{1, \dots, N\}$ and define $Y_\ell = \sum_i Y_{\ell i} \mathbb{1}\{S_\ell = i\}$ as the observed realized test score of student ℓ . A standard approach

¹³We exclude indicators for VRAEM and bilingual (see Figure 1), due to a lack of within-teacher variation in the data.

¹⁴This approach follows recent work in the school/teacher value-added literature (Abdulkadiroğlu et al., 2020; Ahn et al., 2025). This specification for match effects nests other approaches used in the teacher value-added literature, which either assume constant effects (Chetty et al., 2014b) or constant effects within student sub-populations (Biasi et al., 2021; Bates et al., 2025).

in the teacher value-added literature is to assume selection on observables such that

$$\mathbb{E}[Y_\ell | X_{\ell 1}, X_{\ell 2}, S_\ell = i] = \delta_{0i} + X'_{\ell 1} \delta_{1i} + X'_{\ell 2} \delta_2. \quad (8)$$

Under this assumption, and provided that there is sufficient within-teacher variation in $X_{\ell 1}$, a regression of Y_ℓ on teacher fixed effects, $X_{\ell 1}$ interacted with teacher fixed effects, and $X_{\ell 2}$ yields unbiased estimates of $(\delta_{0i}, \delta_{1i}, \delta_2)$. Although value-added methods relying on selection on observables have been shown to yield forecast unbiased estimates of teacher effects in various settings (Chetty et al., 2014a; Bacher-Hicks et al., 2014), this assumption remains contested (Rothstein, 2017).

To relax the selection on observables assumption of standard value-added models, we employ an alternative approach that exploits the data and structure imposed on teacher school choices to recover unbiased estimates of the value-added parameters under weaker assumptions. Following Abdulkadiroğlu et al. (2020), we adapt the control function approach of Dubin and McFadden (1984) to correct for selection on teachers' unobserved preference for their assigned school:

$$\mathbb{E}[Y_\ell | X_{\ell 1}, X_{\ell 2}, S_\ell = i, \eta_{ij(\ell)}] = \delta_{0i} + X'_{\ell 1} \delta_{1i} + X'_{\ell 2} \delta_2 + \varphi \times (\eta_{ij(\ell)} - \mu_\eta), \quad (9)$$

where μ_η is Euler's constant. The parameter φ captures selection on unobserved teacher-school match effects that would make teachers more or less productive in schools that give them larger utility. Note that any selection on other teacher-specific unobserved preferences over schools is captured by δ_{0i} .

Applying the law of iterated expectations, we can write

$$\mathbb{E}[Y_\ell | X_{\ell 1}, X_{\ell 2}, S_\ell = i, D_i, x_i, d_i, s_i] = \delta_{0i} + X'_{\ell 1} \delta_{1i} + X'_{\ell 2} \delta_2 + \varphi \lambda_{j(\ell)}(D_i, x_i, d_i, s_i), \quad (10)$$

where λ_j denotes the control function corresponding to the expected unobserved preference shock for school j conditional on the teacher characteristics (x_i, d_i, s_i) and school j being the chosen alternative of teacher i

$$\lambda_j(D_i, x_i, d_i, s_i) := \mathbb{E}[\eta_{ijt} - \mu_\eta | x_i, d_i, s_i, j = D_{it}]$$

$$= \int \left[\log \left(\sum_{k \in \Omega(s_{it}) \cup \{0, J+1\}} \exp(u_{ikt}) \right) - u_{iD_{it}t} \right] f(\theta_i | x_i, j = D_{it}) d\theta_i.$$

where $f(\theta_i | x_i, j = D_{it})$ denotes the density function of θ_i conditional on x_i and school j being chosen by teacher i .

Under this assumption, OLS yields unbiased estimates of $(\delta_{0i}, \delta_{1i}, \delta_2, \varphi)$ after plugging-in the control function term. Note that we exclude $(d_{ij(\ell)})$ and $(w_{j(\ell)})$ from directly entering students' potential outcomes. This provides important variation necessary to identify φ . Distance to home has been argued to be a valid preference shifter in school choice environments (Walters, 2018; Abdulkadiroğlu et al., 2017). Consistent with this exclusion restriction, columns (1)-(3) of Table A.8 show that, conditional on the controls included in the outcome equation, distance is not systematically related to other predetermined school-level outcomes not included in the achievement equation such as student dropout rates, end-of-year examination failure rates, or year-to-year teacher retention rates. Columns (4)-(6) of Table A.8 report analogous results for wages. We also find no evidence of a direct effect of wage increases on student achievement when the recruitment channel is muted (see Column (2) in Table 3), supporting the validity of this exclusion restriction.¹⁵

The resulting estimates of the teacher specific coefficients $(\hat{\delta}_{0i}, \hat{\delta}_{1i})$ are unbiased but noisy.¹⁶ Recent work has highlighted the benefits of using Empirical Bayes (EB) shrinkage when making decisions involving ranking and selection in the presence of such statistical uncertainty (Gu and Koenker, 2023; Walters, 2024). We denote $\delta_i = (\delta_{0i}, \delta_{1i})$ and impose the following hierarchical structure

$$\hat{\delta}_i \mid \delta_i, \hat{\Omega}_i \sim \mathcal{N}(\delta_i, \hat{\Omega}_i), \quad \delta_i \mid x_i, \hat{\Omega}_i \sim \mathcal{N}(\mu'_\delta x_i, \Sigma_\delta).$$

We estimate $(\mu_\delta, \Sigma_\delta)$ by MLE and construct the estimated posterior means for the teacher value-added coefficients δ_i

¹⁵Note that this approach does not capture selection on unobservables generated by endogenous sorting of students into schools or endogenous classroom formation. Addressing these additional margins would require separate sources of exogenous variation affecting students' school choices or within-school classroom assignments.

¹⁶Accounting for parameter uncertainty in the control function terms through a parametric bootstrap would increase standard errors of $\hat{\delta}_i$ by 0.007% on average (see Appendix B.2). Therefore, we do not adjust the second-stage covariance matrices throughout the rest of the estimation procedure.

$$\hat{\delta}_i^* = (\hat{\Omega}_i^{-1} + \hat{\Sigma}_\delta^{-1})^{-1}(\hat{\Omega}_i^{-1}\hat{\delta}_i + \hat{\Sigma}_\delta^{-1}\hat{\mu}'_\delta x_i).$$

Throughout the rest of the analysis we use $\hat{\delta}_{0i}^*$ as a measure of teacher i 's average effectiveness (ATE) and $X'_{\ell 1}\hat{\delta}_{i1}^*$ as a measure of teacher i 's comparative advantage with students with characteristics $X_{\ell 1}$ (match effect). We refer to the sum of the ATE and the match effect components as the value-added of teacher i for student ℓ .

Finally, we allow the teacher value-added coefficients δ_i to correlate with the random coefficients θ_i . This captures the link between teacher effectiveness and latent factors affecting their preferences over schools that could reveal their intrinsic motivation. It also captures any potential selection on teacher effectiveness that would occur as a result of counterfactual compensation policies. To identify and estimate the covariance matrix $\Sigma_{\theta,\delta}$, we leverage within-teacher variation in the characteristics of the chosen school. Intuitively, when teachers with similar observables choose the outside option or low/high wage positions at a higher frequency, this reveals information about their latent type θ_i . As we have an unbiased estimate of the teacher value-added coefficients $\hat{\delta}_i$, we can then identify and estimate the distribution of θ_i conditional on δ_i , which pins down $\Sigma_{\theta,\delta}$. A formal proof of identification and details on estimation can be found in Appendix B.1.¹⁷

5 Estimation Results

5.1 Teacher Preferences

We report the full set of preference estimates in Tables A.9, A.10, and A.11. Table 4 reports the bottom decile, median, and top decile of the implied distribution of willingness-to-pay, measured in percentage of the monthly base wage, for various non-wage attributes contained in the vectors z_{jt} and d_{ijt} entering (2). Consistent with the survey evidence displayed in Table 1, teachers value local amenities, good teaching conditions, and geographical proximity. The median teacher would be willing to give up 11% of her baseline wage for a one-standard-deviation increase in the amenity/infrastructure quality index. The median willingness to

¹⁷Since value-added is observed only for teachers in the outcome sample, we impose that, conditional on the teacher characteristics entering the distributions of both random coefficients and value added, the relationship between preference heterogeneity and teacher effectiveness captured by $\Sigma_{\theta,\delta}$ is common across both the outcome and choice samples.

Table 4: Willingness to Pay for Non-Pecuniary Job Attributes

	Bottom Decile	Median	Top Decile
	(1)	(2)	(3)
Amenity Index	0.028	0.114	0.450
Multigrade	-1.880	-0.435	-0.075
Single Teacher	-4.263	-1.031	-0.193
Bilingual	-3.270	-0.747	0.338
Share of Female Students	-0.790	-0.103	0.156
Share of Quechua Students	-4.011	-0.898	1.003
Average Math Score	-0.293	0.001	0.236
Average Spanish Score	-0.018	0.015	0.110
Distance Spline: 0-20km	-0.451	-0.124	-0.034
Distance Spline: 20-100km	-0.232	-0.065	-0.018
Distance Spline: 100km+	-0.039	-0.010	-0.003


NOTES. This table displays the bottom decile, median, and top decile of the distribution of the estimated preference parameters across teachers, expressed in fractions of the baseline monthly wage in 2016 (S/ 1,396).

pay to avoid teaching in multi-grade, single-teacher, or bilingual schools ranges between 43% and 103% of the base wage. The median cost of traveling one kilometer further away from home is decreasing with distance and ranges from 12% to 1% of the baseline wage.¹⁸ The median teacher has a preference for schools enrolling relatively lower shares of female and Quechua students. Importantly, Columns (1)-(3) of Table 4 show that these valuations are highly heterogeneous across teachers. Non-pecuniary attributes thus induce substantial vertical and horizontal differentiation between school vacancies offering the same wage.

Schools located in rural areas tend to have worse teaching conditions, worse local amenities, and can potentially be very remote, compared to urban schools. This translates into large spatial differences in the utility of accepting a job. Figure 6 plots cell-average differences in the non-pecuniary components of the indirect utility, measured in monthly wages (in Peruvian Soles, S/), between rural schools and the average urban school, for teachers who score above the 75th percentile in the competency test. The vast majority of cells categorized as extremely rural would need to offer wage bonuses that range from two to four times the actual S/ 500 monthly wage bonus in order to be as attractive as the average urban school.

¹⁸One kilometer measured in geodesic distance may entail substantial travel time in some regions of Perú due to poor road infrastructure quality. This may partly explain the large magnitudes of the estimated moving costs.

Figure 6: Rural vs. Urban Differences in Non-Pecuniary Utility



NOTES. To construct this figure, we restrict the sample to teachers in the upper quartile of the score distribution and simulate the non-pecuniary part of utility, as measured in wages: $\frac{U_{ijt} - \alpha_i w_{jt}}{\alpha_i}$ for each teacher-school pair using the estimated parameters from equation (2) and fixed draws of η_{ijt} and θ_i . Then, we compute the median of the difference in those utilities for each rural school with the average urban school. Finally, we average across schools at the level of equally spaced cells of dimension 100×60 in the population-travel-time space.

Overall, preference estimates suggest that the current wage bonus policy only partially compensates for urban-rural differences in non-pecuniary utility. The current teacher allocation is thus likely to remain highly unequal, favoring urban schools. Figure 6 also shows substantial variation in non-pecuniary utility within geographic areas that receive the same bonus, suggesting that the rural wage bonus may be better targeted across schools and locations. This information, along with the substantial heterogeneity of preferences over non-pecuniary job characteristics (as shown in Table 4), may be leveraged to lower these utility differences and help design better targeted teacher compensation policies, an issue we later return to.

5.2 Teacher Value Added

Table 5 displays our estimates of the distribution of the teacher value-added coefficients δ_i and the coefficient φ capturing selection on unobserved preferences.¹⁹ As shown in Column (2), average effectiveness (ATE) varies substantially across teachers. Moving one standard deviation up the distribution of average effectiveness implies an increase in student score of

¹⁹Column (1) of Table 5 only shows estimates of the mean of δ_i for the baseline category. Estimates of the interaction effects for the conditional mean of δ_i are displayed in Table A.13.

Table 5: Distribution of Value-Added Parameters

	Mean Baseline	SD	Correlation		Covariance RC	
			Lagged Score	Female	Wage	Outside Option
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Math test scores</i>						
ATE	0	0.384 (0.006)	0.078 (0.031)	0.252 (0.082)	-0.049 (0.019)	-0.089 (0.046)
Lagged Score	0.498 (0.011)	0.150 (0.004)		0.021 (0.082)	-0.019 (0.011)	0.001 (0.023)
Female	-0.072 (0.014)	0.097 (0.011)			-0.008 (0.015)	-0.006 (0.030)
Unobserved Preferences (φ)	-0.007 (0.006)	-	-	-	-	-
<i>Panel B: Spanish test scores</i>						
ATE	0	0.329 (0.007)	0.279 (0.029)	-0.035 (0.101)	-0.062 (0.019)	-0.105 (0.043)
Lagged Score	0.549 (0.011)	0.169 (0.004)		-0.150 (0.102)	-0.005 (0.011)	0.026 (0.026)
Female	0.009 (0.014)	0.074 (0.014)			-0.021 (0.015)	-0.025 (0.029)
Unobserved Preferences (φ)	-0.007 (0.006)	-	-	-	-	-

NOTES. This table displays estimates of the mean of the value-added coefficients δ_i for the baseline category and the coefficient φ capturing selection on unobserved preferences through the control function term (Column 1), estimates of the standard deviations of δ_i (Column 2), estimates of the correlation matrix of δ_i (Columns 3 and 4), and estimates of the covariances between the teacher value-added coefficients δ_i and the random coefficients θ_i (Columns 5 and 6). Standard errors are in parentheses.

0.38 σ in math and 0.33 σ in Spanish.

Column (1) shows that the control function model attributes very little variation in student outcomes to match-specific utility that teachers get from their school assignment. As a result, shrunk estimates of the ATE derived from the control function model and the standard value-added model are very close to each other (see Figure A.5). As evidenced in prior work, this suggests that the selection on observables assumption in test score value-added models is reasonable when controlling for lagged outcomes (Chetty et al., 2014a).

Teachers widely differ in their comparative advantage in teaching students with different characteristics. Students with second-grade test scores one standard deviation below average can experience fourth-grade test score gains of 0.15-0.17 σ by being matched with a teacher with similar average effectiveness but with a one standard deviation higher match effect.

Table 6: Variance Decomposition of Teacher Value-Added

	Math	Spanish
	(1)	(2)
Total variance teacher value-added	0.087	0.042
Variance ATE	0.068	0.026
Variance match effects	0.011	0.015
$2 \times \text{Cov}(\text{ATE}, \text{match effects})$	0.007	0.001

NOTES. This table displays the variance of teacher value-added decomposed as the variance of the ATE, the match effects, and the covariance term. Variance components are derived using shrunk estimates of the teacher value-added coefficients.


Teachers who are effective in math, on average, tend to be particularly effective with female students, whereas teachers who are effective in Spanish, on average, tend to be particularly effective with students having higher initial levels of achievement (Columns 3 and 4). This evidence is consistent with findings in [Ahn et al. \(2025\)](#); [Graham et al. \(2023\)](#), which document that students lagging behind have the largest potential gains from teacher reallocation based on comparative advantage. It also points to large potential gains for rural areas from attracting better teachers along *both* absolute and comparative advantages.

Finally, Columns (5) and (6) of Table 5 display the estimated covariance matrix between the random coefficients (θ_i) and the value-added coefficients (δ_i). Teachers who are less responsive to wage differences are more effective, on average, and have a higher comparative advantage with students with lower prior achievement. Teachers with a higher outside option are less effective, on average. These findings play an important role in the counterfactual analysis that we present in Section 6, as they show that teachers attracted by wage increases are likely to be negatively selected. Not accounting for this correlation would bias predictions of teacher-school sorting under counterfactual compensation schedules.

Overall, depending on the subject taught, Table 6 shows that 62-78% of the total variance in teacher value-added can be explained by differences in average effectiveness, while the remaining 22-38% is explained by differences in teachers' comparative advantage. Table A.14 in the Appendix shows that our observed teacher characteristics x_i explain 16-25% of the overall variance in the ATE, but do not explain variation in the match effect coefficients.²⁰ Conditioning on the random coefficients θ_i on top of the observed characteristics x_i slightly lowers the variance of the ATE but does not affect the variance of the match effect coefficients.

²⁰Table A.13 shows that teacher competency scores, experience in the private sector and gender strongly associate with their average effectiveness in math and Spanish.

Figure 7: Model Fit – Sorting and Inequality



NOTES. Panel A compares estimates of the threshold crossing effects displayed in Table 2 with the corresponding simulated threshold-crossing effects estimates averaged over 10 simulated teacher-school matches. To simulate an equilibrium teacher-school match, we take a random draw of type-I EV shocks η_{ijt} , and a random draw of the random coefficients θ_i from its distribution conditional on x_i and $\hat{\delta}_i$. We then construct indirect utilities for every teacher-school pair and run the serial dictatorship algorithm to simulate the assignment mechanism (see Section 2.1). When using priority indices and test scores as outcomes, we exclude simulated matches where the threshold-crossing effect on the probability of filling vacancies is statistically significant at the 95% level to ensure validity of the RD design. The reported confidence intervals have 95% coverage. Panel B shows the urban-rural gap in teacher average effects in math and Spanish, and in teacher competency scores in the data, along with the simulated teacher-school match.

5.3 Model Fit

We assess whether the estimated model replicates the observed teacher sorting patterns through several exercises. To simulate an equilibrium teacher-school match, we take a random draw of type-I extreme value (EV) shocks η_{ijt} , and a random draw of the random coefficients θ_i from its distribution conditional on x_i and $\hat{\delta}_i$. We then construct indirect utilities for every teacher-school pair and run the serial dictatorship algorithm to simulate the assignment mechanism (see Section 2.1).

We first test whether the model can replicate the threshold-crossing effects on teacher sorting discussed in Section 3. This provides a direct assessment of the validity of the wage elasticities implied by our estimates. Panel A of Figure 7 shows that the RD estimates derived from simulated teacher-school matches fall within the 95% confidence interval of the threshold-crossing effects estimated from the data (see Table 2). We then assess whether the model can replicate the overall sorting patterns observed in the data. Panel B of Figure 7 shows that the estimated model accurately replicates the urban-rural gap in teacher value-added and competency scores. Table A.15 in the Appendix further shows that model predictions closely match a wide range of additional moments of the distribution of matched

teacher and school characteristics.

We further leverage our RD design to assess the robustness of our estimates of teacher effectiveness. Using the same sample as that in Column (1) of Table 3, we replicate our RD analysis on residualized student test scores constructed by differencing out the teacher value-added components (see Table A.16 in Appendix). Estimated threshold-crossing effects are still positive, but substantially muted and statistically insignificant. As teacher effectiveness is the only input in the student achievement production function (7) that varies non-smoothly at the population cutoff, these results provide additional support for the validity of our value-added estimates.

6 Counterfactuals

6.1 Gains from Teacher Reallocation and Recruitment

We first use our teacher value-added estimates to study how teacher sorting contributes to the observed spatial gaps in student outcomes. Column (1) of Table 7 reports the average teacher effectiveness in urban and rural schools in the status quo scenario. Despite the large effects around the population threshold documented in the RD analysis of Section 3, teacher sorting remains highly unequal. Teachers working in rural schools have 0.14σ lower value-added than in urban schools. Match effects are uniformly unimportant, indicating that teachers do not sort based on their comparative advantages in the status quo (see Panel B of Table 7).

We next characterize the potential gains from teacher reallocation. We restrict attention to the teacher-classroom matches in our outcome sample. For each recruitment year $t \in \{2016, 2018\}$, let $\mathcal{I}_t$ denote the set of teachers and $\mathcal{C}_t$ the set of classrooms observed in these matches. Let $\mathcal{L}_{ct}$ denote the set of students enrolled in classroom c and $n_{ct} = |\mathcal{L}_{ct}|$ its class size. Let a_{ict} be an indicator equal to one if teacher i is assigned to classroom c in year t . For each recruitment year, we solve the following linear assignment problem:

$$\begin{aligned} \max_{\{a_{ict}\}} \quad & \sum_{c \in \mathcal{C}_t} \sum_{i \in \mathcal{I}_t} a_{ict} \left[\frac{1}{n_{ct}} \sum_{\ell \in \mathcal{L}_{ct}} \left(\widehat{\delta}_{0i}^* + X'_{\ell 1} \widehat{\delta}_{i1}^* \right) \right] \\ \text{s.t.} \quad & \sum_{i \in \mathcal{I}_t} a_{ict} = 1, \quad \forall c \in \mathcal{C}_t, \quad \sum_{c \in \mathcal{C}_t} a_{ict} = 1, \quad \forall i \in \mathcal{I}_t, \quad a_{ict} \in \{0, 1\}. \end{aligned} \tag{11}$$

The first set of constraints ensures that each classroom is assigned exactly one teacher, while

Table 7: Teacher Sorting and Inequality

	Data	Value-added Maximization
	(1)	(2)
<i>Panel A: ATE Math</i>		
Urban	0.112	0.053
Rural	-0.029	0.050
<i>Other Rural</i>	0.001	0.056
<i>Extremely Rural</i>	-0.085	0.038
Total	0.052	0.052
<i>Panel B: Match Effects Math</i>		
Urban	0.004	0.044
Rural	0.003	0.091
<i>Other Rural</i>	0.002	0.075
<i>Extremely Rural</i>	0.005	0.122
Total	0.003	0.064

NOTES. This table summarizes the spatial distribution of teacher effectiveness decomposed into average effectiveness (Panel A) and match effects (Panel B) in the teacher-school match observed in the data (Column 1) and in the teacher value-added maximization counterfactual (Column 2).

the second ensures that each teacher is assigned to exactly one classroom. The objective maximizes the teacher-specific component of predicted student achievement. The remaining components of the student achievement production function do not vary with the identity of the teacher assigned to a given classroom and therefore do not affect the solution. We divide the individual contributions of each student in the objective function by class size to avoid mechanically penalizing rural schools that enroll fewer students on average.

This exercise provides an estimate of the gains attainable by reallocating the existing pool of public teachers. Because neither the pool of teachers nor the set of classrooms changes, the counterfactual can generate gains only by changing who teaches whom. Aggregate gains can thus arise only by exploiting teachers' comparative advantages with students having different characteristics. We solve the problem separately for the 2016 and 2018 recruitment drives and pool the resulting teacher-classroom assignments when reporting the counterfactual outcomes in Column (2) of Table 7.

We find that reallocating the pool of existing teachers would increase student achievement in rural schools by 0.17σ ($=0.08\sigma$ in ATE + 0.09σ in Match Effects) and moderately decrease value-added in urban schools by 0.02σ . This would translate into aggregate gains in teacher value-added of 0.06σ . Importantly, we find that the gains in aggregate achievement are

accompanied by a substantial reduction in spatial inequality. Because the pool of teachers is held fixed, the aggregate ATE is unchanged by construction. The optimal assignment nevertheless redistributes average teacher effectiveness across locations: ATE increases by 0.08σ in rural schools and decreases by 0.06σ in urban schools. This result suggests that equity and efficiency are not sharply opposed in this setting.

We contrast this reallocation benchmark with an exercise that changes the composition of the public-teacher pool. For each recruitment year t , let $\mathcal{O}_t$ denote the set of applicants who chose the outside option in that year but are observed teaching in the other year. We can therefore estimate their average effectiveness even though they did not enter public teaching through the centralized recruitment drive in year t . We augment the set of assigned teachers $\mathcal{I}_t$ with these applicants and, holding fixed the number of filled teaching positions, select the teachers with the highest ATE. As in the reallocation exercise, we solve this problem separately for 2016 and 2018 and pool the resulting outcomes across recruitment years. Unlike the reallocation counterfactual, this exercise does not seek to exploit school-specific match effects. Instead, it isolates the gains from replacing teachers with relatively low average effectiveness with more effective applicants who otherwise choose the outside option. We find potential aggregate gains of up to 0.11σ , compared with the 0.06σ gain from reallocating the existing teacher pool. This suggests that large efficiency and equity gains could be attained by leveraging this extensive margin rather than by reallocating the existing pool of teachers to exploit match effects.

6.2 Counterfactual Compensation Policy

We seek to design a wage-setting protocol that systematically leverages information on teacher effectiveness and preferences over schools to achieve the efficiency and equity gains previously highlighted. We assume that we have two policy levers: (i) the priority index used to rank teachers in the centralized assignment mechanism, and (ii) the wage w_j offered by each school j . As in the existing compensation system, all vacancies within a school offer the same wage.

As schools can open more than one vacancy, we define the policy objective as ensuring that each school j in a predefined set $\mathcal{S} \subseteq \{1, \dots, J\}$ is assigned at least one teacher whose effectiveness for that school exceeds a threshold $\bar{c}$. We denote by Y_{ij} teacher i 's predicted value-added at school j , which incorporates both average effectiveness and the match effect associated with the students enrolled at that school. For any wage vector $w = (w_1, \dots, w_J)$, let

$D(w) = \{D_i(w)\}_{i=1}^N$ denote the associated teacher-school assignment where teachers would rank schools according to the indirect utilities specified in Equation (2). We seek to construct a counterfactual wage schedule satisfying

$$\max_{\{i:D_i(w)=j\}} Y_{ij} \geq \bar{c}, \quad \forall j \in \mathcal{S}, \quad (\text{C1})$$

and

$$D(w) \text{ is stable given } w \text{ and school priorities } Y_{ij}. \quad (\text{C2})$$

Condition (C1) imposes that our objective is attained under the resulting teacher-school match $D(w)$. Condition (C2) imposes that the teacher-school match $D(w)$ is implementable using Deferred Acceptance if we set wages to w and priorities to Y_{ij} .

Mirroring the wage-proposal dynamics of generalized Deferred Acceptance in matching-with-contracts environments (Hatfield and Milgrom, 2005), we construct the counterfactual schedule through iterative wage adjustments. Starting from the initial wage vector, we compute the stable teacher-school assignment at the prevailing wages using Y_{ij} as school priorities.²¹ For each school in $\mathcal{S}$ that is not assigned a teacher with $Y_{ij} \geq \bar{c}$, we increase its wage by one point on a discrete wage grid and recompute the entire assignment. We repeat this process until every targeted school satisfies (C1). Recomputing the assignment after each wage adjustment accounts for the equilibrium effects of changes in compensation on teacher sorting across all schools.

The resulting allocation satisfies the two constraints by construction. Condition (C2) holds at each iteration because the assignment is obtained by running Deferred Acceptance at the current wage vector. At termination, Condition (C1) holds for every targeted school by the stopping rule. The procedure also converges in a finite number of iterations provided that the number of teachers clearing the threshold $\bar{c}$ is sufficiently high. Because wages are increased only at schools that fail to attract a teacher above the effectiveness threshold at the prevailing assignment, each adjustment directly contributes to attaining the policy objective.

We apply the ascending-wage procedure described above using teacher value-added in math as the effectiveness measure Y_{ij} , using a step size for wage increases of $S/20$, and defining $\mathcal{S}$ as the set of rural schools with at least one filled vacancy under the status quo.

²¹To preserve strategy-proofness of the assignment mechanism, we use teacher-proposing Deferred Acceptance to derive the assignment.

Table 8: Counterfactual Teacher Compensation Policy


	Status Quo	Counterfactual Policy
	(1)	(2)
<i>Panel A: ATE Math</i>		
Urban	0.112	0.138
Rural	-0.029	0.089
<i>Other Rural</i>	0.001	0.093
<i>Extremely Rural</i>	-0.085	0.079
Total	0.052	0.119
<i>Panel B: Match Effects Math</i>		
Urban	0.004	0.005
Rural	0.003	0.009
<i>Other Rural</i>	0.002	0.005
<i>Extremely Rural</i>	0.005	0.022
Total	0.003	0.007
<i>Panel C: Monthly Wage Bonus</i>		
Urban	0	0
Rural	235	227
<i>Other Rural</i>	86	153
<i>Extremely Rural</i>	500	365
Total	122	120

NOTES. This table summarizes the distribution of teacher average effectiveness (Panel A) and match effects (Panel B) under the simulated equilibrium teacher-school matches resulting from different wage bonus policies. Column (1) corresponds to the simulated match in the status quo while Column (2) corresponds to the simulated match under the wage-setting protocol described in Section 6.2 when setting Y_{ij} as math test score value-added and setting the threshold $\bar{c}$ such that the policy is budget neutral relative to the status quo policy. To simulate the equilibrium teacher-school match in each scenario, we construct indirect utilities for every teacher-school pair by taking the same random draw of type-I EV shocks η_{ijt} and random coefficients θ_i from their distribution conditional on x_i and $\hat{\delta}_i$.

We implement the procedure for different values of $\bar{c}$ and select a counterfactual wage schedule that is approximately budget neutral relative to the status quo policy. We report the distribution of teacher value-added under this counterfactual compensation policy in Column (2) of Table 8.

Our results show that large equity and efficiency gains are attainable under alternative wage schedules. Teacher value added increases by 0.12σ in rural schools and by 0.03σ in urban schools. In contrast with the optimal reallocation counterfactual shown in Column (2) of Table 7, teacher value-added increases in the aggregate mostly through an increase in average effectiveness (ATE), i.e. by attracting into the assigned-teacher pool higher-quality applicants who chose the outside option under the status quo. This suggests that inducing

Figure 8: Distribution of Counterfactual Wage Bonuses Relative to Status Quo



NOTES. This figure plots cell-averaged differences between the monthly wage bonuses attributed to each school in the counterfactual policy described in Column 2 of Table 8 and the status quo policy. Cells are equally spaced with dimension 100×60 in the population-travel-time (in minutes) space.

teacher reallocation to exploit match effects might be less cost-effective than leveraging the extensive margin of recruitment by attracting applicants with higher average effectiveness into public teaching.

We further characterize how our counterfactual wage bonuses are distributed across schools. Panel C of Table 8 shows that the counterfactual policy reduces the overall weight put on extremely rural schools and redistributes it partly to other rural schools. Figure 8 shows a more disaggregated version of these figures by plotting cell-average differences in the population-travel-time space between our counterfactual wages and the status quo policy. This uncovers substantial heterogeneity within the extremely rural category. Schools located close to the 120 travel-time threshold are substantially over-subsidized in the status quo policy, while the most remote schools are under-subsidized. Taken together, this evidence highlights the benefits of acquiring information on teachers' preferences and effectiveness to design more effective teacher compensation policies.

7 Conclusion

This paper establishes that the design of teacher compensation can largely exacerbate or alleviate structural inequalities in schooling outcomes. We assemble rich data on applicants

and jobs posted in the nationwide centralized recruitment drive for public teachers in Perú. Wage rigidity induces teachers to sort on non-pecuniary aspects of employment, resulting in school choices that are skewed towards urban areas. We leverage a policy reform that increased teacher compensation in rural schools and show that it was effective at attracting higher quality teachers and improving student learning.

To go beyond the local estimated effects of the policy and understand the potential equity and efficiency gains that can be achieved through alternative teacher compensation schemes, we build and estimate a model of teacher sorting across schools and student achievement. Teachers are heterogeneous in their preferences over wage and non-wage attributes, which induces vertical and horizontal differentiation across jobs. Teacher sorting maps into student achievement through a teacher value-added model in which teacher effectiveness is heterogeneous and interacts with students' prior achievement and gender.

Using these estimates, we characterize the gains available through changes in both the allocation and composition of the public-teacher workforce and study whether they can be attained through alternative compensation policies. Holding fixed the existing pool of teachers and classrooms, the assignment that maximizes teacher effectiveness across classrooms would increase student achievement in rural schools by 0.17σ while reducing it in urban schools by only 0.02σ , generating an aggregate gain of 0.06σ . An alternative benchmark that expands the teacher pool to include applicants who chose the outside option and selects those with the highest average effectiveness generates potential aggregate gains of up to 0.11σ . We then develop a counterfactual wage-setting procedure that uses information on teachers' preferences and effectiveness to pursue these gains while accounting for equilibrium sorting through the centralized assignment mechanism. We find that a substantial reduction in the urban-rural gap in teacher effectiveness can be achieved at approximately the same cost as the status quo policy. Unlike the unconstrained reallocation benchmark, these gains arise primarily from attracting applicants with higher average effectiveness into public teaching rather than from inducing existing teachers to sort according to their comparative advantages. This suggests that leveraging the extensive margin of recruitment may be more cost-effective than inducing teacher reallocation to exploit match effects.

Beyond the specific setting of our analysis, our findings suggest that incorporating measures of preferences and productivity in the design of worker compensation is a promising alternative to rigid wage schedules or market-based wage setting in a variety of other labor

markets. Aligning the private and social returns of worker-firm matches through informed compensation design can result in more equitable and efficient allocations. We believe that this approach is increasingly relevant from a policy perspective, given the widespread availability of administrative data on centralized labor markets as well as recent developments in the tools that enable researchers and practitioners to leverage such data to infer the preferences of participating agents ([Agarwal and Budish, 2021](#)).

References

- Abdulkadiroğlu, Atila, Nikhil Agarwal, and Parag A Pathak**, “The Welfare Effects of Coordinated Assignment: Evidence from the New York City High School Match,” *American Economic Review*, 2017, *107* (12), 3635–89.
- , **Parag A Pathak, Jonathan Schellenberg, and Christopher R Walters**, “Do Parents Value School Effectiveness?,” *American Economic Review*, 2020, *110* (5), 1502–1539.
- Agarwal, Nikhil**, “Policy Analysis in Matching Markets,” *American Economic Association P&P*, 2017, *107* (5), 246–50.
- and **Eric Budish**, “Market Design,” in “Handbook of Industrial Organization,” Vol. 5, Elsevier, 2021, pp. 1–79.
- Ahn, Tom, Esteban M Aucejo, and Jonathan James**, “The Importance of Student-Teacher Matching: A Multidimensional Value-Added Approach,” *Review of Economics and Statistics*, 2025, pp. 1–45.
- Araujo, M Caridad, Pedro Carneiro, Yyannú Cruz-Aguayo, and Norbert Schady**, “Teacher quality and learning outcomes in kindergarten,” *The Quarterly Journal of Economics*, 2016, *131* (3), 1415–1453.
- Bacher-Hicks, Andrew, Thomas J Kane, and Douglas O Staiger**, “Validating teacher effect estimates using changes in teacher assignments in Los Angeles,” Technical Report, National Bureau of Economic Research 2014.

- Banerjee, Abhijit and Lakshmi Iyer**, “History, institutions, and economic performance: The legacy of colonial land tenure systems in India,” *American Economic Review*, 2005, 95 (4), 1190–1213.
- Barrera-Osorio, Felipe and Dhushyanth Raju**, “Teacher performance pay: Experimental evidence from Pakistan,” *Journal of Public Economics*, 2017, 148 (C), 75–91.
- Bates, Michael, Michael Dinerstein, Andrew C Johnston, and Isaac Sorkin**, “Teacher Labor Market Policy and the Theory of the Second Best,” *The Quarterly Journal of Economics*, 2025, 140 (2), 1417–1469.
- Bau, Natalie and Jishnu Das**, “Teacher value added in a low-income country,” *American Economic Journal: Economic Policy*, 2020, 12 (1), 62–96.
- Berry, Steven, James Levinsohn, and Ariel Pakes**, “Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market,” *Journal of Political Economy*, 2004, 112 (1), 68–105.
- Bertoni, Eleonora, Gregory Elacqua, Luana Marotta, Matías Martínez, Humberto Santos, and Sammara Soares**, “Is School Funding Unequal in Latin America? A Cross-Country Analysis. CEPA Working Paper No. 20-11.,” *Stanford Center for Education Policy Analysis*, 2020.
- Biasi, Barbara**, “The Labor Market for Teachers under Different Pay Schemes,” *American Economic Journal: Economic Policy*, August 2021, 13 (3), 63–102.
- , **Chao Fu, and John Stromme**, “Equilibrium in the Market for Public School Teachers: District Wage Strategies and Teacher Comparative Advantage,” *NBER Working Paper 28530*, 2021.
- Boyd, Donald, Hamilton Lankford, Susanna Loeb, and James Wyckoff**, “Analyzing the Determinants of the Matching of Public School Teachers to Jobs: Disentangling the Preferences of Teachers and Employers,” *Journal of Labor Economics*, 2013, 31 (1), 83–117.
- Brown, Christina and Tahir Andrabi**, “Inducing positive sorting through performance pay: Experimental evidence from Pakistani schools,” *University of California at Berkeley Working Paper*, 2020.

- Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik**, “Robust Nonparametric Confidence Intervals for Regression Discontinuity Designs,” *Econometrica*, November 2014, *82*, 2295–2326.
- , —, and —, “Optimal Data-Driven Regression Discontinuity Plots,” *Journal of the American Statistical Association*, December 2015, *110* (512), 1753–1769.
- Cattaneo, Matias D., Michael Jansson, and Xinwei Ma**, “Simple Local Polynomial Density Estimators,” *Journal of the American Statistical Association*, 2020, *115* (531), 1449–1455.
- Chetty, Raj, John Friedman, and Jonah Rockoff**, “Measuring the Impact of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood,” *American Economic Review*, 2014, *104*(9), 2633–2679.
- , —, and —, “Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates,” *American Economic Review*, 2014, *104*(9), 2593–2632.
- Clotfelter, Charles, Elizabeth Glennie, Helen Ladd, and Jacob Vigdor**, “Would Higher Salaries Keep Teachers in High-Poverty Schools? Evidence from a Policy Intervention in North Carolina,” *Journal of Public Economics*, 2008, *92* (5), 1352–1370.
- Clotfelter, Charles T, Helen F Ladd, and Jacob Vigdor**, “Who teaches whom? Race and the distribution of novice teachers,” *Economics of Education Review*, 2005, *24* (4), 377–392.
- de Ree, Joppe, Karthik Muralidharan, Menno Pradhan, and Halsey Rogers**, “Double for Nothing? Experimental Evidence on an Unconditional Teacher Salary Increase in Indonesia,” *The Quarterly Journal of Economics*, 2018, *133* (2), 993–1039.
- Dell, Melissa**, “The persistent effects of Peru’s mining mita,” *Econometrica*, 2010, *78* (6), 1863–1903.
- Dubin, Jeffrey A and Daniel L McFadden**, “An econometric analysis of residential electric appliance holdings and consumption,” *Econometrica*, 1984, pp. 345–362.

- Fack, Gabrielle, Julien Grenet, and Yinghua He**, “Beyond Truth-Telling: Preference Estimation with Centralized School Choice and College Admissions,” *American Economic Review*, 2019, *109* (4), 1486–1529.
- Finan, Frederico, Benjamin A. Olken, and Rohini Pande**, “The Personnel Economics of the State,” in “Handbook of Field Experiments,” Vol. 2, North Holland, 2017.
- Fox, Jeremy T, Kyoo il Kim, Stephen P Ryan, and Patrick Bajari**, “The random coefficients logit model is identified,” *Journal of Econometrics*, 2012, *166* (2), 204–212.
- Gerard, François, Miikka Rokkanen, and Christoph Rothe**, “Bounds on treatment effects in regression discontinuity designs with a manipulated running variable,” *Quantitative Economics*, 2020, *11* (3), 839–870.
- Gershenson, Seth, Cassandra MD Hart, Joshua Hyman, Constance A Lindsay, and Nicholas W Papageorge**, “The long-run impacts of same-race teachers,” *American Economic Journal: Economic Policy*, 2022, *14* (4), 300–342.
- Gilligan, Daniel O., Naureen Karachiwalla, Ibrahim Kasirye, Adrienne M. Lucas, and Derek Neal**, “Educator Incentives and Educational Triage in Rural Primary Schools,” *Journal of Human Resources*, 2022, *57* (1), 79–111.
- Graham, Bryan S, Geert Ridder, Petra Thiemann, and Gema Zamarro**, “Teacher-to-classroom assignment and student achievement,” *Journal of Business & Economic Statistics*, 2023, *41* (4), 1328–1340.
- Gu, Jiaying and Roger Koenker**, “Invidious comparisons: Ranking and selection as compound decisions,” *Econometrica*, 2023, *91* (1), 1–41.
- Hatfield, John William and Paul R. Milgrom**, “Matching with Contracts,” *American Economic Review*, September 2005, *95* (4), 913–935.
- INEI**, “Encuesta Nacional de Hogares (ENAHO),” <https://www.ilo.org/surveyLib/index.php/catalog/6803/study-description> Instituto Nacional de Estadística e Informática 2016.
- Jackson, C. Kirabo**, “What Do Test Scores Miss? The Importance of Teacher Effects on Non-Test Score Outcomes,” *Journal of Political Economy*, 2018, *126* (5), 2072–2107.

- Judd, Kenneth L**, *Numerical Methods in Economics*, MIT press, 1998.
- Leaver, Clare, Owen Ozier, Pieter Serneels, and Andrew Zeitlin**, “Recruitment, Effort, and Retention Effects of Performance Contracts for Civil Servants: Experimental Evidence from Rwandan Primary Schools,” *American Economic Review*, July 2021, 111 (7), 2213–46.
- Mansfield, Richard K**, “Teacher quality and student inequality,” *Journal of Labor Economics*, 2015, 33 (3), 751–788.
- Menzel, Konrad**, “Large Matching Markets as Two-Sided Demand Systems,” *Econometrica*, 2015, 83 (3), 897–941.
- MINEDU**, “Encuesta Nacional a Docentes de Instituciones Educativas Estatales y No Estatales (ENDO),” <https://hdl.handle.net/20.500.12799/4872> Ministerio de Educacion 2014.
- Muralidharan, Karthik and Venkatesh Sundararaman**, “Teacher Performance Pay: Experimental Evidence from India,” *Journal of Political Economy*, 2011, 119 (1), 39–77.
- Rivkin, Steven G, Eric A Hanushek, and John F Kain**, “Teachers, Schools, and Academic Achievement,” *Econometrica*, 2005, 73 (2), 417–458.
- Rosen, Sherwin**, “The theory of equalizing differences,” *Handbook of Labor Economics*, 1986, 1, 641–692.
- Roth, Alvin E and Marilda Sotomayor**, “Two-Sided Matching,” *Handbook of Game Theory with Economic Applications*, 1992, 1, 485–541.
- Rothstein, Jesse**, “Teacher Quality Policy When Supply Matters,” *American Economic Review*, January 2015, 105 (1), 100–130.
- , “Measuring the impacts of teachers: Comment,” *American Economic Review*, 2017, 107 (6), 1656–1684.
- Sokoloff, Kenneth L and Stanley L Engerman**, “History lessons: institutions, factor endowments, and paths of development in the new world,” *Journal of Economic Perspectives*, 2000, 14 (3), 217–232.

- Tincani, Michela M**, “Teacher Labor Markets, School Vouchers, and Student Cognitive Achievement: Evidence from Chile,” *Quantitative Economics*, 2021, *12* (1), 173–216.
- Walters, Christopher R**, “The demand for effective charter schools,” *Journal of Political Economy*, 2018, *126* (6), 2179–2223.
- , “Empirical Bayes methods in labor economics,” in “Handbook of Labor Economics,” Vol. 5, Elsevier, 2024, pp. 183–260.
- World Bank**, “World Development Report 2018 : Learning to Realize Education’s Promise,” Technical Report, Washington, DC: World Bank 2018.

Appendices

A Additional Tables and Figures

Table A.1: Descriptive Statistics for Analysis Samples

	Mean	SD
<i>Panel A: Applicant characteristics (choice sample)</i>		
Age	37.734	(6.922)
Female	0.742	(0.438)
Quechua	0.212	(0.409)
Experience in the public sector (≥ 3 yrs.)	0.573	(0.495)
Experience in the private sector (≥ 1 yr.)	0.479	(0.500)
Degree	0.297	(0.457)
Overall competency score (std)	-0.062	(1.013)
N. of teachers	36,698	
N. of applications	50,319	
<i>Panel B: Vacancy characteristics (choice sample)</i>		
Monthly Wage (soles)	1,821.671	(386.627)
Single Teacher	0.007	(0.081)
Multigrade	0.192	(0.394)
Urban	0.531	(0.499)
Extremely Rural	0.154	(0.361)
Other Rural	0.314	(0.464)
Bilingual	0.101	(0.301)
Border	0.062	(0.240)
VRAEM	0.032	(0.176)
Time-to-travel to provincial capital	120.951	(270.289)
Locality population (log)	9.104	(3.632)
Amenity/infrastructure index	0.109	(0.970)
N. of school	7,826	
N. of vacancies	22,743	
<i>Panel C: Student characteristics (outcome sample)</i>		
Age	9.202	(0.666)
Female	0.496	(0.500)
Quechua	0.122	(0.328)
Spanish test scores (std)	-0.097	(0.988)
Math test scores (std)	-0.051	(1.001)
N. of teachers	4,834	
N. of students	96,508	

NOTES. This table reports summary statistics for the characteristics of applicants (Panel A) and vacancies (Panel B) that were part of the assignment mechanism for temporary positions in 2016 and 2018 and constitute our choice sample. Panel C reports the characteristics of 4th-grade students in our outcome sample.

Table A.2: School and Locality Characteristics by Rurality

	Rural		Urban	
	Mean	SD	Mean	SD
<i>Panel A: School characteristics</i>				
Single Teacher	0.0142	(0.118)	0	(0)
Multigrade	0.404	(0.491)	0.00579	(0.0759)
Bilingual	0.193	(0.395)	0.0196	(0.139)
Border	0.0701	(0.255)	0.0541	(0.226)
VRAEM	0.0462	(0.210)	0.0194	(0.138)
<i>Panel B: Locality amenities</i>				
Amenity/infrastructure index	-0.385	(0.932)	0.544	(0.774)
Poverty index	0.964	(1.424)	0.00995	(0.887)
Electricity	0.909	(0.287)	0.997	(0.0509)
Drinking water	0.674	(0.469)	0.941	(0.236)
Sewage	0.429	(0.495)	0.925	(0.263)
Water tower	0.256	(0.436)	0.597	(0.491)
Medical clinic	0.638	(0.480)	0.916	(0.277)
Meal center	0.251	(0.434)	0.584	(0.493)
Community phone	0.0910	(0.288)	0.122	(0.327)
Internet access point	0.181	(0.385)	0.883	(0.321)
Bank	0.0587	(0.235)	0.628	(0.483)
Public library	0.0389	(0.193)	0.488	(0.500)
Police	0.199	(0.399)	0.598	(0.490)
<i>Panel C: School infrastructure</i>				
Teachers room	0.134	(0.341)	0.324	(0.468)
Sport pitch	0.167	(0.373)	0.325	(0.468)
Courtyard	0.174	(0.379)	0.376	(0.484)
Auditorium	0.0868	(0.282)	0.195	(0.396)
Administrative office	0.484	(0.500)	0.666	(0.472)
First-aid room	0.0108	(0.103)	0.0777	(0.268)
Computer lab	0.371	(0.483)	0.834	(0.372)
Workshop	0.0490	(0.216)	0.304	(0.460)
Science lab	0.0766	(0.266)	0.420	(0.494)
Library	0.433	(0.495)	0.721	(0.448)
Personal computer	0.785	(0.411)	0.916	(0.278)

NOTES. This table reports the average and standard deviation of various school characteristics by rurality. The sample includes all vacancies in the choice sample.

Table A.3: Survey Respondents Characteristics

	All applicants (2016) (1)	Survey respondents (2)
Age	34.546 (6.184)	35.201 (6.546)
Female	0.725 (0.447)	0.723 (0.448)
Quechua	0.109 (0.312)	0.084 (0.277)
Experience in the public sector (≥ 3 yrs.)	0.621 (0.485)	0.614 (0.487)
Experience in the private sector (≥ 1 yr.)	0.674 (0.469)	0.709 (0.454)
Degree	0.559 (0.497)	0.563 (0.496)
Overall competency score (std)	1.593 (0.390)	1.608 (0.457)
N. of teachers	22,784	5,550

NOTES. This table reports average teacher characteristics for the sample of survey respondents (Column 2) and the sample of applicants in the 2016 recruitment drive for permanent teaching positions (Column 1). Standard deviations are reported in parentheses.

Table A.4: Covariate Smoothness around the Population Cutoff

	Choice Sample			Outcome Sample		
	Mean (BW)	RD estimate		Mean (BW)	RD estimate	
	(1)	(2)	(3)	(4)	(5)	(6)
<i>School characteristics</i>						
Distance to the provincial capital	163.021	5.468	(43.241)	139.613	30.572	(36.632)
Distance >120 min.	0.384	-0.004	(0.068)	0.378	0.207	(0.123)
Other wage bonuses (S/.)	72.187	0.643	(14.152)	53.482	16.806	(22.907)
Single Teacher	0.000	0.000	(0.002)	0.008	0.004	(0.008)
Multigrade	0.312	-0.075	(0.075)	0.162	0.128	(0.117)
Bilingual	0.205	0.064	(0.051)	0.175	-0.089	(0.070)
Border	0.057	-0.116	(0.056)	0.050	-0.100	(0.058)
VRAEM	0.045	0.054	(0.031)	0.042	-0.010	(0.051)
Number of students	114.153	0.322	(6.342)	115.502	-5.331	(10.140)
Indigenous language students	0.268	-0.021	(0.064)	0.209	0.011	(0.082)
% indigenous language students	0.171	-0.019	(0.048)	0.138	-0.095	(0.068)
% proficient students (math)	9.739	1.547	(3.053)	10.648	3.788	(4.530)
% proficient students (spanish)	13.152	1.075	(3.520)	15.189	2.356	(3.826)
<i>Village amenities</i>						
Amenity/infrastructure index	-0.219	0.029	(0.124)	-0.119	-0.038	(0.148)
Poverty index	0.684	0.353	(0.217)	0.733	0.330	(0.319)
Electricity	0.960	0.068	(0.047)	0.964	0.058	(0.037)
Drinking water	0.756	0.008	(0.074)	0.784	0.076	(0.116)
Sewage	0.506	0.047	(0.080)	0.531	0.013	(0.112)
Water tower	0.246	-0.044	(0.075)	0.374	-0.251	(0.144)
Medical clinic	0.779	0.042	(0.066)	0.767	0.110	(0.120)
Meal center	0.262	0.109	(0.069)	0.315	-0.014	(0.100)
Community phone	0.078	0.018	(0.044)	0.068	0.037	(0.055)
Internet access point	0.144	0.004	(0.050)	0.183	-0.044	(0.082)
Bank	0.036	-0.001	(0.027)	0.035	0.001	(0.035)
Public library	0.023	-0.023	(0.023)	0.025	-0.018	(0.028)
Police	0.194	0.017	(0.061)	0.178	-0.178	(0.109)
<i>School amenities</i>						
Teachers room	0.184	-0.044	(0.052)	0.221	-0.136	(0.078)
Sport pitch	0.201	-0.037	(0.065)	0.218	-0.066	(0.117)
Courtyard	0.182	-0.106	(0.074)	0.189	-0.027	(0.090)
Auditorium	0.102	-0.005	(0.044)	0.076	0.053	(0.092)
Administrative office	0.530	0.024	(0.084)	0.534	-0.034	(0.117)
First-aid room	0.010	0.005	(0.003)	0.017	-0.058	(0.044)
Computer lab	0.436	0.067	(0.082)	0.470	-0.056	(0.126)
Workshop	0.071	-0.027	(0.030)	0.047	-0.000	(0.022)
Science lab	0.081	0.074	(0.041)	0.116	0.043	(0.073)
Library	0.494	0.013	(0.083)	0.494	-0.086	(0.139)
At least a personal computer	0.793	0.132	(0.064)	0.781	0.171	(0.073)

NOTES. This table examines whether schools located just above or just below the population threshold differ in village and school amenities (measured in 2013). Columns (1)–(3) focus on the choice sample, while columns (4)–(6) focus on the outcome sample. The table reports the mean of each variable for control schools (i.e., schools located to the right of the population cutoff within the chosen bandwidth) and the robust bias-corrected regression-discontinuity estimates, together with their standard errors, obtained using the estimator proposed by [Calonico et al. \(2014\)](#).

Table A.5: Decomposition of wage increases around the population cutoff

	(1)	(2)	(3)
	Overall	Time < 120	Time > 120
Bonus	244.205 (26.533)	26.503 (12.824)	409.121 (18.497)
Mean dep. var. (Lower Bonus)	1844.006	1800.000	1897.416
Bandwidth	143.908	142.384	240.045
Schools	1081	671	956
Observations	2600	1301	2297

NOTES. This table reports the effect of crossing the population threshold on teacher wages. In all columns, the dependent variable is defined as the sum of the base wage and the wage bonuses described in Figure 1 (for schools that satisfy the criteria). The sample used is the choice sample. Column (2) only considers schools that do not comply with the travel time cutoff (closer than 120 minutes). Column (3) considers schools that comply with the travel time cutoff, that is, farther than 120 minutes from the provincial capital. In all columns, the sample excludes urban schools (i.e., those located in areas with more than 2,000 inhabitants) and schools situated within 30 minutes of the nearest provincial capital. Cells report bias-corrected regression-discontinuity estimates obtained using the robust estimator proposed by Calonico et al. (2014). Regressions are estimated within a mean-squared-error-optimal bandwidth, reported at the bottom of the table. Standard errors are clustered at the school×year level.

Table A.6: Sharp RD Bounds Under Potential Manipulation

	(1)	(2)
	Preferences	Teacher Score (std.)
Upper bound	0.142	0.503
Lower bound	0.083	0.275
CI	[0.047 - 0.177]	[0.150 - 0.628]
Bandwidth	162.27	193.38
Schools	1157	1388
Observations	2575	3068

NOTES. This table reports the RD bounds (Gerard et al., 2020) for the threshold crossing effect on two outcomes that are subject to potential selection due to the assignment of applicants to vacancies. In Column (1) the outcome variable is the rank in which a vacancy was chosen in the assignment mechanism (normalized so that it takes values from zero to one); in Column (2) it is the standardized competency score obtained by the teachers in the centralized test. The 95% confidence intervals are obtained through 1000 bootstrap replications. The sample used is the choice sample, further restricted to vacancies filled in the national recruitment drive. Additionally, the sample excludes urban schools (i.e., those located in areas with more than 2,000 inhabitants) and schools situated within 30 minutes of the nearest provincial capital.

Table A.7: Teacher Score (Std.)–Difference-in-Discontinuities

	(1)	(2)	(3)
	2016	2018	Pooled
Post-policy	-0.022 (0.068)	0.172 (0.063)	0.069 (0.055)
Post-policy $\times$ High Bonus	0.203 (0.085)	0.274 (0.081)	0.243 (0.070)
Bandwidth	134.340	216.811	159.777
Schools	930	2003	1498
Observations	3373	6120	6758

NOTES. This table reports the effect of crossing the population threshold on the standardized competency score of contract teachers across different school years. Post-policy is an indicator equal to one in the post-treatment period—that is, 2016 in Column (1), 2018 in Column (2), and both years in Column (3)—and zero in the pre-policy year (2015). Post-policy $\times$ High Bonus is the interaction between the former and an indicator for municipalities with a population below 500 inhabitants, which, from 2016 onward are eligible for the higher wage bonus. All regressions include school fixed effects. Cells report conventional regression-discontinuity estimates obtained using a triangular kernel function and the optimal bandwidth selected from a standard RD specification without interaction terms, estimated on the corresponding sample. Standard errors are clustered at the school level.

Table A.8: Balance Tests Instruments

	(1)	(2)	(3)	(4)	(5)	(6)
	Dropout	Failing rate	Teacher retention	Dropout	Failing rate	Teacher retention
Intercept	0.023 (0.003)	0.050 (0.004)	0.874 (0.020)	0.022 (0.005)	0.043 (0.006)	0.887 (0.028)
Distance (1,000 km)	0.0007 (0.0035)	0.0025 (0.0046)	-0.0215 (0.0212)	-	-	-
Wage (1,000 soles)	-	-	-	0.0005 (0.0017)	0.0035 (0.0022)	-0.0082 (0.0104)
School Lagged Math Score	0.001 (0.002)	0.005 (0.003)	0.014 (0.015)	0.001 (0.002)	0.005 (0.003)	0.015 (0.015)
School Lagged Spanish Score	-0.010 (0.003)	-0.012 (0.003)	0.002 (0.016)	-0.010 (0.003)	-0.011 (0.003)	0.001 (0.016)
School Lagged Math Score ²	0.001 (0.001)	-0.005 (0.002)	-0.005 (0.010)	0.001 (0.001)	-0.005 (0.002)	-0.004 (0.010)
School Lagged Spanish Score ²	-0.001 (0.001)	0.003 (0.002)	0.005 (0.008)	-0.001 (0.001)	0.003 (0.002)	0.005 (0.008)
School Female	-0.012 (0.004)	0.007 (0.006)	0.036 (0.028)	-0.012 (0.004)	0.007 (0.006)	0.037 (0.028)
School Quechua	-0.009 (0.002)	0.005 (0.002)	-0.063 (0.011)	-0.009 (0.002)	0.005 (0.002)	-0.063 (0.011)
School Age	0.026 (0.002)	0.095 (0.003)	-0.048 (0.015)	0.026 (0.002)	0.095 (0.003)	-0.047 (0.015)
Amenity Index	-0.004 (0.001)	-0.001 (0.001)	0.003 (0.004)	-0.004 (0.001)	-0.001 (0.001)	0.003 (0.004)
Border	-0.003 (0.002)	0.000 (0.003)	0.024 (0.014)	-0.003 (0.002)	0.000 (0.003)	0.025 (0.014)
Single Teacher	0.026 (0.004)	0.008 (0.005)	0.067 (0.026)	0.026 (0.004)	0.008 (0.005)	0.069 (0.026)
Multigrade	-0.001 (0.002)	0.004 (0.002)	-0.042 (0.011)	-0.001 (0.002)	0.004 (0.002)	-0.041 (0.011)
Travel Time	0.003 (0.001)	0.000 (0.002)	-0.036 (0.008)	0.003 (0.001)	-0.000 (0.002)	-0.035 (0.008)
log(Population)	0.000 (0.000)	-0.001 (0.000)	0.002 (0.001)	0.000 (0.000)	-0.000 (0.000)	0.002 (0.001)
Travel Time $\times$ log(Pop)	-0.001 (0.000)	0.000 (0.000)	0.003 (0.001)	-0.001 (0.000)	0.000 (0.000)	0.002 (0.001)

NOTES. This table reports school–teacher-level regressions of student dropout rates, end-of-year examination failure rates, and year-to-year teacher retention on distance between the teacher’s residence and the school (Columns (1)-(3)) and on the wage offered by the school (Columns (4)-(6)). All specifications include school-level averages of lagged math and Spanish achievement and their squares, student composition, school and locality characteristics, log population, and the interaction between log population and travel time. Standard errors are reported in parentheses.

Table A.9: Estimates of Teacher Preferences – μ_θ

		× Female	× Teacher Score	× Quechua	× Exp Public > 3	× Exp Private > 0
	(1)	(2)	(3)	(4)	(5)	(6)
Wage	-0.394 (0.130)	-0.237 (0.049)	0.176 (0.030)	-0.154 (0.060)	-0.028 (0.046)	-0.059 (0.048)
Outside Option ($j = 0$)	2.895 (0.450)	0.306 (0.264)	-1.250 (0.139)	-0.385 (0.312)	-1.448 (0.258)	0.647 (0.263)

NOTES. This table displays the estimates of the parameters of the conditional mean of θ_i . Standard errors are in parentheses.

Table A.10: Estimates of Teacher Preferences – Σ_θ

	Wage	Outside Option ($j = 0$)
	(1)	(2)
Wage	0.986 (0.073)	0.647 (0.095)
Outside Option ($j = 0$)		2.036 (0.052)

NOTES. This table displays the estimates of the variance covariance matrix of θ_i . The diagonal elements are the standard deviations of θ_i while the off diagonal elements display $\text{corr}(\theta_i)$. Standard errors are in parentheses.

Table A.11: Estimates of Teacher Preferences – (Θ, γ)

	$\times$ Female	$\times$ Teacher Score	$\times$ Quechua	$\times$ Exp Public > 3	$\times$ Exp Private > 0	
	(1)	(2)	(3)	(4)	(5)	(6)
Amenity Index	0.043 (0.027)	0.014 (0.022)	-0.015 (0.011)	-0.002 (0.030)	0.025 (0.023)	0.034 (0.024)
Border	0.155 (0.100)	-0.194 (0.085)	0.084 (0.049)	-0.156 (0.135)	0.017 (0.086)	-0.072 (0.089)
VRAEM	-0.121 (0.176)	-0.285 (0.115)	0.200 (0.072)	-0.147 (0.132)	0.292 (0.129)	-0.009 (0.122)
Multigrade	-0.250 (0.068)	-0.234 (0.055)	0.026 (0.032)	-0.131 (0.065)	0.190 (0.058)	0.027 (0.058)
Single Teacher	-0.192 (0.239)	-0.631 (0.199)	-0.147 (0.129)	-0.064 (0.304)	-0.229 (0.223)	-0.048 (0.245)
Bilingual	-1.042 (0.105)	0.113 (0.077)	-0.155 (0.046)	0.972 (0.082)	0.199 (0.083)	0.056 (0.082)
Time	-0.225 (0.039)	-0.046 (0.028)	-0.148 (0.017)	0.040 (0.044)	0.080 (0.038)	-0.136 (0.041)
Time ²	-0.002 (0.000)	0.002 (0.000)	0.001 (0.000)	0.000 (0.000)	0.000 (0.000)	-0.003 (0.000)
log(Population)	0.026 (0.061)	0.045 (0.048)	0.039 (0.024)	-0.093 (0.057)	-0.026 (0.045)	-0.004 (0.046)
log(Population) ²	-0.016 (0.003)	-0.000 (0.002)	0.001 (0.001)	0.002 (0.003)	0.001 (0.002)	0.003 (0.002)
Time $\times$ log(Pop)	0.047 (0.006)	-0.002 (0.004)	0.010 (0.002)	-0.006 (0.007)	-0.013 (0.005)	0.029 (0.006)
School Lagged Math Score	0.019 (0.072)	-0.016 (0.061)	0.081 (0.034)	0.106 (0.071)	-0.076 (0.063)	0.049 (0.063)
School Lagged Spanish Score	-0.006 (0.077)	0.018 (0.064)	-0.009 (0.036)	-0.020 (0.076)	-0.006 (0.068)	0.028 (0.068)
School Female	-0.174 (0.150)	0.019 (0.125)	0.081 (0.063)	0.171 (0.149)	0.184 (0.125)	-0.145 (0.126)
School Quechua	-0.884 (0.120)	0.081 (0.084)	-0.020 (0.050)	1.525 (0.095)	-0.158 (0.087)	-0.125 (0.086)
School Age	0.080 (0.077)	-0.080 (0.061)	-0.092 (0.034)	-0.053 (0.078)	-0.005 (0.068)	-0.002 (0.068)
Previous School	4.508 (0.070)	0.048 (0.055)	-0.259 (0.027)	-0.696 (0.073)	-0.256 (0.056)	0.748 (0.053)
Distance Spline: $[0, 20)$	-0.095 (0.004)	-0.001 (0.003)	-0.025 (0.002)	-0.013 (0.004)	0.006 (0.003)	-0.001 (0.003)
Distance Spline: $[20, 100)$	-0.048 (0.001)	0.000 (0.001)	-0.006 (0.000)	0.002 (0.001)	-0.003 (0.001)	0.003 (0.001)
Distance Spline: ≥ 100	-0.007 (0.000)	0.001 (0.000)	0.001 (0.000)	0.000 (0.000)	-0.001 (0.000)	-0.001 (0.000)
Other School ($j = J + 1$)	2.122 (0.447)	0.472 (0.259)	-1.184 (0.135)	-0.398 (0.307)	0.122 (0.252)	0.287 (0.258)

NOTES. This table displays our estimates of the parameters Θ_1 , Θ_2 , γ_2 . Standard errors are in parentheses.

Table A.12: Estimates of Teacher Value-Added – δ_2

	Math	Spanish
	(1)	(2)
Lagged Score ²	-0.001 (0.003)	-0.011 (0.003)
Quechua	-0.047 (0.018)	-0.023 (0.018)
Age	-0.132 (0.006)	-0.095 (0.006)
School Lagged Score	0.306 (0.056)	0.422 (0.061)
School Lagged Score ²	0.056 (0.051)	0.198 (0.048)
School Female	-0.228 (0.265)	-0.725 (0.268)
School Quechua	0.268 (0.299)	0.250 (0.303)
School Age	-0.759 (0.140)	-0.419 (0.144)
Classroom Lagged Score	-0.162 (0.036)	-0.203 (0.039)
Classroom Lagged Score ²	-0.038 (0.030)	-0.019 (0.027)
Classroom Female	0.655 (0.150)	0.600 (0.152)
Classroom Quechua	-0.349 (0.312)	-0.236 (0.312)
Classroom Age	0.076 (0.080)	0.031 (0.078)
Class Size	0.020 (0.008)	0.010 (0.008)
Class Size ²	-0.000 (0.000)	-0.000 (0.000)
Amenity Index	0.042 (0.021)	0.045 (0.022)
Border	0.017 (0.137)	0.005 (0.141)
Multigrade	0.198 (0.090)	0.140 (0.091)
Single Teacher	0.083 (0.229)	-0.280 (0.235)
Travel Time	0.118 (0.080)	0.056 (0.081)
log(Pop)	-0.019 (0.012)	-0.029 (0.012)
Travel Time $\times$ log(Pop)	-0.015 (0.011)	-0.004 (0.011)
Control Function Missing	-0.039 (0.046)	-0.012 (0.046)

NOTES. This table displays the estimates of the parameters associated with school and classroom characteristics in the student achievement production function. When the control-function term is missing, we set it to zero and include an indicator for a missing control-function term.

Table A.13: Estimates of Teacher Value-Added – μ_δ

	$\times$ Female	$\times$ Quechua	$\times$ Teacher Score	$\times$ Exp Public > 3	$\times$ Exp Private > 0
	(1)	(2)	(3)	(4)	(5)
<i>Panel A: Math</i>					
ATE	0.074 (0.017)	0.058 (0.022)	0.107 (0.009)	-0.003 (0.017)	0.096 (0.017)
Lagged Score	0.007 (0.009)	-0.013 (0.011)	0.027 (0.004)	-0.011 (0.008)	-0.001 (0.009)
Female	0.004 (0.012)	-0.002 (0.015)	-0.002 (0.006)	0.003 (0.011)	-0.014 (0.011)
<i>Panel B: Spanish</i>					
ATE	0.125 (0.016)	-0.065 (0.021)	0.077 (0.008)	-0.015 (0.015)	0.098 (0.016)
Lagged Score	0.018 (0.010)	-0.011 (0.012)	0.020 (0.005)	-0.017 (0.009)	0.009 (0.009)
Female	-0.010 (0.012)	-0.011 (0.015)	0.006 (0.006)	0.009 (0.011)	-0.018 (0.011)

NOTES. This table displays the estimates of the parameters of the interaction effects in the conditional mean of the distribution of the teacher value-added coefficients δ_i . Standard errors are in parentheses.

Table A.14: Teacher Characteristics and Teacher Value Added

	$\text{Var}(\cdot \theta_i, x_i)$	$\text{Var}(\cdot x_i)$	$\text{Var}(\cdot)$
	(1)	(2)	(3)
<i>Panel A: Math</i>			
ATE	0.112	0.116	0.139
Lagged Score	0.022	0.023	0.024
Female	0.009	0.009	0.009
<i>Panel B: Spanish</i>			
ATE	0.043	0.047	0.063
Lagged Score	0.028	0.029	0.030
Female	0.005	0.005	0.006

NOTES. This table displays the variance of the value-added coefficients conditional on x_i and θ_i (Column 1), conditional on x_i only (Column 2) and unconditional (Column 3). Variance components are derived using the shrunk value-added estimates.

Table A.15: Model Fit – Distribution of Matched Teacher-School Characteristics

	Mean		SD	
	Data	Model	Data	Model
	(1)	(2)	(3)	(4)
Wage (in thousands of soles)	1.876	1.880	0.382	0.387
Amenity Index	0.057	0.049	0.967	0.968
Border	0.060	0.060	0.237	0.237
VRAEM	0.029	0.031	0.169	0.172
Multigrade	0.222	0.224	0.415	0.417
Single Teacher	0.007	0.007	0.085	0.085
Bilingual	0.108	0.108	0.311	0.310
Travel Time (in hours)	1.970	1.978	3.989	3.894
log(Population)	8.783	8.789	3.628	3.628
School Lagged Math Score	-0.034	-0.035	0.683	0.683
School Lagged Spanish Score	-0.185	-0.187	0.697	0.701
School Female	0.494	0.493	0.141	0.144
School Quechua	0.153	0.154	0.329	0.331
School Age	0.064	0.063	0.326	0.335
Distance	69.387	72.183	155.083	124.889
Previous School	0.140	0.133	0.347	0.340
Other Public School	0.305	0.310	0.460	0.463
Outside Option	0.351	0.357	0.477	0.479


NOTES. This table displays moments of the distribution of matched teacher and school characteristics in the data and in a simulated teacher-school status quo match. To simulate an equilibrium teacher-school match, we take a random draw of type-I EV shocks η_{ijt} , and a random draw of the random coefficients θ_i from its distribution conditional on x_i and $\hat{\delta}_i$. We then construct indirect utilities for every teacher-school pair and run the serial dictatorship algorithm to simulate the assignment mechanism (see Section 2.1)

Table A.16: Model Fit – Threshold Crossing Effects on Residualized Test Scores

	(1)	(2)
	Math	Spanish
High Bonus	0.187	0.121
	(0.134)	(0.136)
Bandwidth	122.514	134.728
Observations	3547	3843


NOTES. This table reports regression-discontinuity estimates of the effect of crossing the population threshold on students' test scores in math and Spanish, after residualizing outcomes by differencing out the teacher value-added components (see Equation (7)). The sample does not consider schools closer than 30 minutes from the provincial capital and schools located in urban areas (population > 2,000 inhabitants). Cells report the bias-corrected regression-discontinuity estimates obtained using the robust estimator proposed in [Calónico et al. \(2014\)](#). Regressions are defined within a mean-square error optimal bandwidth, reported at the bottom part of the table. Standard errors are clustered at the school×year level.

Figure A.1: Correlation between Rurality, Teacher Competency, and Student Achievement




NOTES: This figure illustrates the relationship between teacher competency, student achievement, and different dimensions of rurality. Rurality is proxied using two measures: the logarithm of the population of the village in which the school is located (Panels A and C) and travel time from the school to the nearest provincial capital (Panels B and D). Teacher competency (Panels A and B) is measured by the average score obtained by teachers in the 2016 and 2018 centralized assignment processes, while student achievement (Panels C and D) is measured by the average student score on the math module. Panels E and F show the relationship between teacher competency and student achievement in math and Spanish, respectively. Each circle represents the mean value within a percentile bin (ventile).


Figure A.2: Densities of the Running Variables: Choice Sample




a. Population (2016)



b. Population (2018)




c. Travel Time (2016)




d. Travel Time (2018)

NOTES. This figure displays the empirical densities with the corresponding confidence intervals for two running variables (population and travel time) for both rounds of the teacher recruitment drive (2016 and 2018). 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 considers all schools in the choice sample.


Figure A.3: Densities of the Running Variables: Outcome Sample




a. Population (2016)



b. Population (2018)




c. Travel Time (2016)



d. Travel Time (2018)


NOTES. This figure displays the empirical densities with the corresponding confidence intervals for two running variables (population and travel time) for both rounds of the teacher recruitment drive (2016 and 2018). 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 considers all schools in the outcome sample.

Figure A.4: Rural Bonus and the Origin of Newly Recruited Teachers



NOTES. This figure plots regression-discontinuity estimates of the effect of crossing the population threshold on the probability that newly recruited contract teachers come from schools located in different population bins. Each dot corresponds to an RD coefficient from a regression in which the outcome is an indicator equal to one if the teacher was employed in 2015 in a public school located in a village whose population falls within the bin shown on the x-axis, and zero otherwise. The bins are mutually exclusive; applicants who were employed in 2015 in the same school to which they are assigned in the 2016 and/or 2018 recruitment drive, or who were not employed in any public school in 2015, take the value zero in all regressions. The sample consists of vacancies in the choice sample that were filled through the centralized assignment process. Additionally, the sample excludes urban schools (i.e., those located in areas with more than 2,000 inhabitants) and schools situated within 30 minutes of the nearest provincial capital. The figure reports bias-corrected RD estimates with 95% confidence intervals, computed using the robust estimator of [Calonico et al. \(2014\)](#). Dashed vertical lines indicate the population bins that lie within the optimal bandwidth used in Column (3) of Table 2.

Figure A.5: Control Function and OLS Estimates



NOTES. This figure plots shrunk estimates of the ATE coming from OLS estimates not controlling for selection (x-axis) against the control-function estimates (y-axis) for math test scores (Panel A) and Spanish test scores (Panel B).

B Technical Appendix

B.1 Identification and Estimation of $\Sigma_{\delta,\theta}$

We can write the probability of observing the matching history $\{D_{it}\}_{t=1}^T$ conditional on observed teacher and school characteristics, teachers' choice sets and their value-added coefficients δ_i :

$$\begin{aligned} \mathbb{P}(\{D_{it}\}_{t=1}^T | x_i, s_i, d_i, \Omega(s_{it}), \delta_i) &= \int \prod_{t=1}^T \frac{\exp\{u_{iD_{it}t}\}}{\exp\{u_{i0t}\} + \exp\{u_{iJ+1t}\} + \sum_{k \in \Omega(s_{it})} \exp\{u_{ikt}\}} \\ &\quad \times \phi(\theta_i | \mu'_\theta x_i + \Sigma_{\theta,\delta} \Sigma_\delta^{-1} (\delta_i - \mu'_\delta x_i), \Sigma_\theta - \Sigma_{\theta,\delta} \Sigma_\delta^{-1} \Sigma_{\delta,\theta}) d\theta_i \end{aligned}$$

[Fox et al. \(2012\)](#) show that this relationship can be inverted to identify the mean and variance of the mixing distribution. The following function $\tilde{\mu}$ is thus identified:

$$\tilde{\mu}(x_i, \delta_i) = \mu'_\theta x_i + \Sigma_{\theta,\delta} \Sigma_\delta^{-1} (\delta_i - \mu'_\delta x_i)$$

From there, we can use variation in δ_i to identify $\Sigma_{\theta,\delta} \Sigma_\delta^{-1}$. Conditional on knowing Σ_δ , we can recover $\Sigma_{\theta,\delta}$.

We estimate $\Sigma_{\theta,\delta}$ by maximizing the joint log-likelihood of observing the matching history $(D_{it})_{t=1}^T$ and the value-added estimates $\hat{\delta}_i$ for $i = 1, \dots, N$ after plugging-in estimates of the preference parameters $(\Theta_1, \Theta_2, \gamma_1, \gamma_2, \beta_1, \kappa)$, the distribution of the random coefficients $(\mu_\theta, \Sigma_\theta)$ and the distribution of value-added coefficients $(\mu_\delta, \Sigma_\delta)$. The log-likelihood has the following expression

$$\begin{aligned} \mathcal{L}(\Sigma_{\theta,\delta}) &= \sum_{i=1}^N \log \int \prod_{t=1}^T \frac{\exp\{u_{iD_{it}t}\}}{\exp\{u_{i0t}\} + \exp\{u_{iJ+1t}\} + \sum_{k \in \Omega(s_{it})} \exp\{u_{ikt}\}} \\ &\quad \times \phi(\hat{\delta}_i | \mu'_\delta x_i + \Sigma_{\delta,\theta} \Sigma_\theta^{-1} (\theta_i - \mu'_\theta x_i), \Sigma_\delta - \Sigma_{\delta,\theta} \Sigma_\theta^{-1} \Sigma_{\theta,\delta} + \hat{\Omega}_i) \phi(\theta_i | \mu'_\theta x_i, \Sigma_\theta) d\theta_i, \end{aligned}$$

where we approximate the integral using 50 draws from a Halton sequence ([Judd, 1998](#)).

B.2 Quantifying Generated-regressor Uncertainty

To approximate how estimation error in the generated control-function terms would affect our second-stage value-added estimates, we implement a parametric bootstrap based on the

estimated asymptotic distribution of the teacher-preference parameters. For each bootstrap draw b , we draw a parameter vector $\Theta^{(b)}$, recompute the corresponding control-function terms $\lambda^{(b)}$, and construct the induced change in the fitted selection correction,

$$r^{(b)} = \hat{\varphi} \left(\lambda^{(b)} - \hat{\lambda} \right).$$

We then compute the change in value-added estimates $\Delta \hat{\delta}_i^{(b)}$ for all i by regressing $-r^{(b)}$ on the same teacher fixed effects, interacted fixed effects, and common covariates used in the outcome equation. Conditional on holding $\hat{\varphi}$ fixed, this calculation is exact and avoids re-estimating the full value-added model for every bootstrap draw. The adjusted covariance matrix for $\hat{\delta}_i$ is then $\hat{\Omega}_i + \text{Var}(\Delta \hat{\delta}_i^{(b)})$ assuming orthogonality between second-stage estimation error and the first-stage-induced perturbation. For 50 bootstrap draws, the implied percentage increase in the standard errors of the teacher value-added estimates has an average of 0.007% and a standard deviation of 0.018% across teachers.