Annenberg Institute
BROWN UNIVERSITY

EdWorkingPaper No. 24-1104

What Happens When We Pay Our Teachers More? Evidence from New Jersey Public Schools

Prasiddha Shakya
Colorado State University

This paper examines the impact of increasing teacher salaries on student outcomes by exploiting variation from the "50K The First Day" campaign that established a $50K salary floor for new teachers across New Jersey school districts. Using school-level data from 2003 to 2019, we employ a staggered difference-in-differences (DiD) approach and first show that the campaign raised salaries for all teachers in New Jersey by approximately $1.5K. Our results indicate that districts implementing the salary increase experienced improvements in 4th grade and high school Math and English Language Arts (ELA) proficiency scores. We also observe modest gains in graduation rate and college enrollment. Analyzing the mechanisms through which these positive effects could have been observed, we rule out teacher migration as a key driver suggesting that the observed improvements are more likely due to changes in teacher motivation and the quality of new teachers entering the profession. Lifting teacher salaries for all teachers—regardless of their performance level—seems to be improving student outcomes in New Jersey.

VERSION: December 2024

Suggested citation: Shakya, Prasiddha. (2024). What Happens When We Pay Our Teachers More? Evidence from New Jersey Public Schools. (EdWorkingPaper: 24 -1104). Retrieved from Annenberg Institute at Brown University: https://doi.org/10.26300/cc21-rg70

What Happens When We Pay Our Teachers More? Evidence from New Jersey Public Schools

Prasiddha Shakya*

December 3, 2024

Link to the Latest Draft

Abstract

This paper examines the impact of increasing teacher salaries on student outcomes by exploiting variation from the "50K The First Day" campaign that established a $50K salary floor for new teachers across New Jersey school districts. Using school-level data from 2003 to 2019, we employ a staggered difference-in-differences (DiD) approach and first show that the campaign raised salaries for all teachers in New Jersey by approximately $1.5K. Our results indicate that districts implementing the salary increase experienced improvements in 4^{th} grade and high school Math and English Language Arts (ELA) proficiency scores. We also observe modest gains in graduation rate and college enrollment. Analyzing the mechanisms through which these positive effects could have been observed, we rule out teacher migration as a key driver suggesting that the observed improvements are more likely due to changes in teacher motivation and the quality of new teachers entering the profession. Lifting teacher salaries for all teachers—regardless of their performance level—seems to be improving student outcomes in New Jersey.

*Department of Economics, Colorado State University, 1771 Campus Delivery, Fort Collins CO 80523-1771, USA. Email: pshakya@colostate.edu
Acknowledgments: I would like to thank my advisors Ray Miller and Sammy Zahran for their insightful feedback and suggestions. Special thanks to Jason Hecht for encouraging me to pursue this area of research during my final year of undergraduate.

1 Introduction

Schoolteachers in the United States are one of the most frequently discussed occupational groups, and the question of whether they are justifiably paid remains a focal point of debate. While existing research has consistently examined how factors like teacher turnover, experience, and principal quality affect student performance,[1] the question of whether higher teacher salaries lead to improved student outcomes remains unanswered.

This paper addresses this question by analyzing the effects of across-the-board salary increases for teachers on student performance in New Jersey public schools. Specifically, we focus on the "50K The First Day" campaign, which set a $50K salary floor for new teachers across school districts. We exploit variation in the timing and intensity of salary changes to identify the causal impact of increasing teacher salaries on student outcomes. Our findings suggest that increasing teacher salaries for all teachers, regardless of their performance level, leads to significant improvements in 4th grade and high school Math and English Language Arts (ELA) proficiency scores, with additional modest gains in graduation rates and college enrollment.

This paper makes three main contributions to the literature. First, to our knowledge, it is among the first studies to examine the causal effect of increasing teacher salaries without allowing districts the discretion to target specific teachers. Unlike previous studies, which focus on salary increases tied to teacher performance,[2] our study looks at a more comprehensive salary increase that affects all teachers equally, regardless of their individual performance. Other studies that answer whether higher pay for all teachers results in better student outcomes have mostly been associative. Our research question is similar to that of Han and Garcia (2022), who used state fixed effects and multilevel mixed-effects models to demonstrate that higher base salaries are linked to significantly improved test scores in mathematics and English Language Arts (ELA). While addressing the same question, our study advances the literature by employing a staggered Differences-in-Differences (DiD) approach, thereby offering a more robust causal analysis.

Second, we explore one key mechanism through which these salary increases may have affected student outcomes: teacher mobility. We find no evidence that salary changes led to substantial teacher migration across school districts. Instead, the improvements seem to stem from higher-quality new hires and increased motivation among current staff. This mechanism contrasts with findings from studies like Baron (2018) and Biasi (2021), who examined Wisconsin's Act 10—a law that weakened unions and introduced performance-based pay. In Wisconsin, salary cuts led to declines in student achievement, and performance-based pay led to sorting of higher-quality teachers into wealthier districts, exacerbating inequality. By studying union-driven salary increases in New Jersey, we provide a unique perspective on how positive compensation changes affect student outcomes without the distortions seen in performance-based pay reforms.[3] This broader lens not only fills an important gap in the

[1] See Eberts and Stone (1988), Miller (2013), Ladd and Sorensen (2017), Henry and Redding (2020), and Ng (2024) for relevant literature.

[2] See Fryer (2013), Goodman and Turner (2013), and Biasi (2021) for performance-based pay studies in the US, and Muralidharan and Sundararaman (2011), Muralidharan et al. (2016), and Hanushek et al. (2019) for studies outside the US context.

[3] Appendix A.1 briefly outlines the theoretical model exploring this mechanism.

literature but also provides new insights into how uniform salary increases can influence educational outcomes at scale without relying on performance-based distinctions that may inherently disadvantage certain educators.

Lastly, we provide a cost-benefit analysis of increasing salaries for all teachers by considering the broader impact on student performance. Our estimations are framed in the context of other meta-study research, such as Jackson et al. (2016) and Jackson and Mackevicius (2024), which demonstrate that increased school spending improves student outcomes through factors like reduced class sizes, increased instructional time, and higher teacher salaries that help attract and retain more qualified educators.[4] These results are in contrast to previous studies by Hanushek (1997, 2003, 2015), which argue that evidence does not robustly support the idea that there is a strong or consistent relationship between student performance and school resources, at least after variations in family inputs are taken into account. To address concerns about endogeneity—where simultaneous increases in spending on resources and teacher salaries might confound results—we demonstrate that school spending on other areas remained relatively constant, allowing us to isolate the effect of salary increases on student outcomes.

The remainder of the paper is organized as follows. Section 2 discusses the "50K The First Day" campaign in detail. After outlining the data (Section 3), the estimation strategy (Section 4), and descriptive statistics (Section 5), Section 6 shows how this campaign increased salary differentials between "treated" and "not yet treated" districts.[5] Given our findings (Sections 7 and 8) that student outcomes improve with higher teacher salaries, Section 9 compares the magnitude of estimated impacts with other comparable studies. Section 10 explores the mechanisms through which salary increases could be influencing these outcomes.

2 Background: "50K the First Day Campaign"

In New Jersey, every school district uses a salary guide with steps that dictate teacher compensation. Table 2 shows an example salary guide where a teacher at Step 1 (no prior experience) with a Bachelor's degree will start at $45K annually. If they hold a Master's degree, they will start at $46K, and with a Master's plus an additional 30 credits, their starting salary will be $46.5K. While these guides serve as references, actual rates can vary across districts. The steps are not tied to specific subjects, though districts struggling to hire for in-demand subjects (e.g., Math) may offer higher step placements to attract candidates.

Step	Bachelors	Masters	Masters + 30 Additional Credits
1	45,000	46,000	46,500
2	45,500	46,500	47,500
3	46,000	47,000	48,500

Table 1: Example Salary Guide

[4] Jackson and Mackevicius (2024) found that increasing spending by $1,000 per pupil over four years improves test scores by 0.0316 standard deviations and raises college-going rates by 2.8 percentage points.

[5] School districts are labeled "treated" if they passed the negotiation that year, while "not yet treated" refers to districts that have yet to implement the salary schedule.

Our paper exploits changes in the salary structure brought about by a bargaining negotiation led by the New Jersey Education Association (NJEA). NJEA is the leading labor union for NJ public school teachers, working to advance the rights, benefits, and interests of its members while promoting quality public education.[6] One of its notable campaigns was "$50K The First Day," which aimed to set a minimum salary of $50K for new teachers at Step 1 of the salary schedule.[7] Implementing such a guide would require all schools in a district to pay Step 1 teachers at least $50K. Figure 1 shows the rollout of the campaign from 2010 to 2015. Minimum salaries negotiated by different school districts varied from $50K to a little under $54K (Figure 2). We account for this variation in treatment intensity in our empirical model.

[6] After the U.S. Supreme Court's Janus v. AFSCME decision in 2018, teaching staff are no longer required to be members of the NJEA. NJEA continues to represent teachers across all districts, with little evidence suggesting that it advocates for improved working conditions only in specific districts.

[7] The NJEA is currently advocating for a "$60K The First Day" campaign, urging districts to raise the starting salary to $60K.

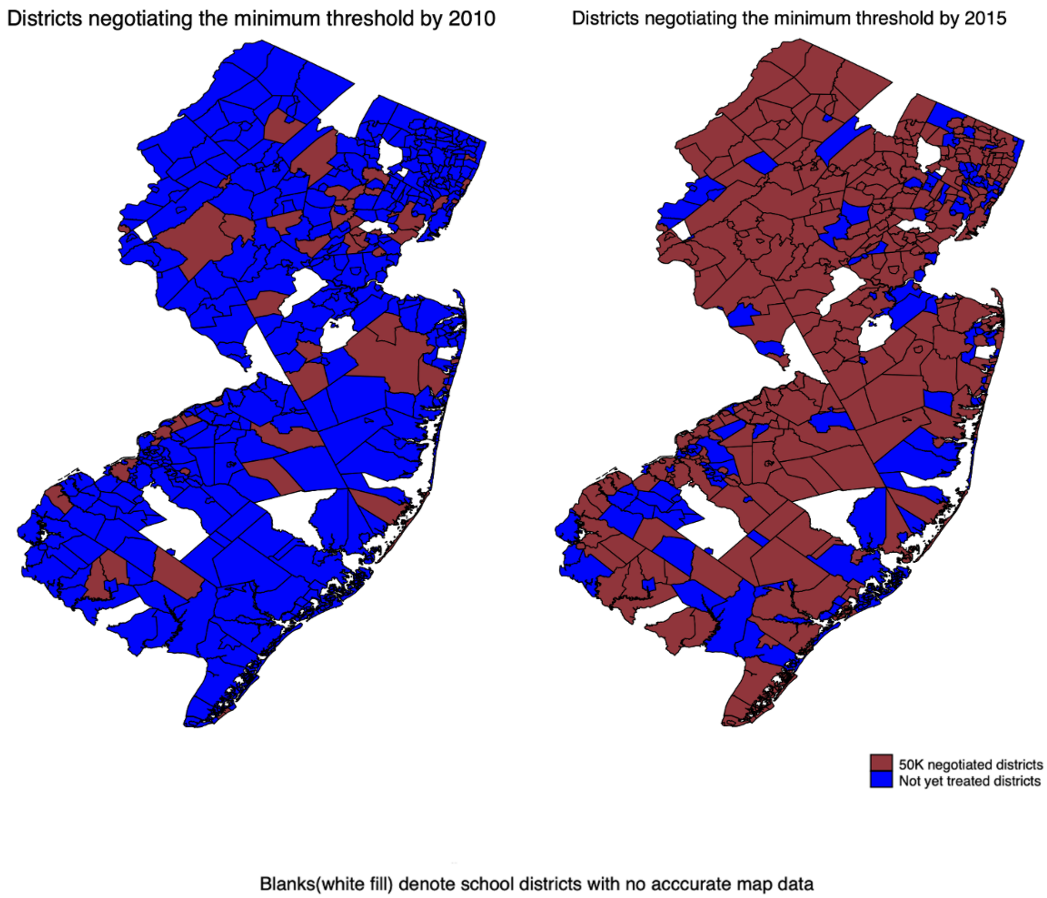

Figure 1:
Number of districts passing the negotiation over the years.

Note: Figure shows the years each school district negotiated the 50K salary schedule. Negotiation always occurs at the school district level. School districts with a Step 1 Salary of at least 50K are colored brown. We exploit the staggered nature of this campaign to estimate the impact of salaries on student outcomes.

New Jersey public schools implement salary guides at the district level, and the natural question to ask is why some districts choose to adopt new schedules while others do not. The most common assumption is that districts refrain from increasing salaries due to budget constraints. NJEA, however, argues that school districts often do have the funds to reallocate toward salary increases. Instead, the decision to adjust salary guides is usually political and driven by concerns over voter reactions, potential backlash from reallocating funds away from veteran teachers toward newer hires, or reluctance to approve settlement percentages above the average. Negotiations are particularly challenging when the policy is perceived as "taking money from the top," i.e., reducing veteran teacher pay.[8] Rather than relying on expanded budgets, many districts negotiated higher salaries by reallocating breakage money—savings

[8] We thank Crystal Inman of NJEA for providing background and context.

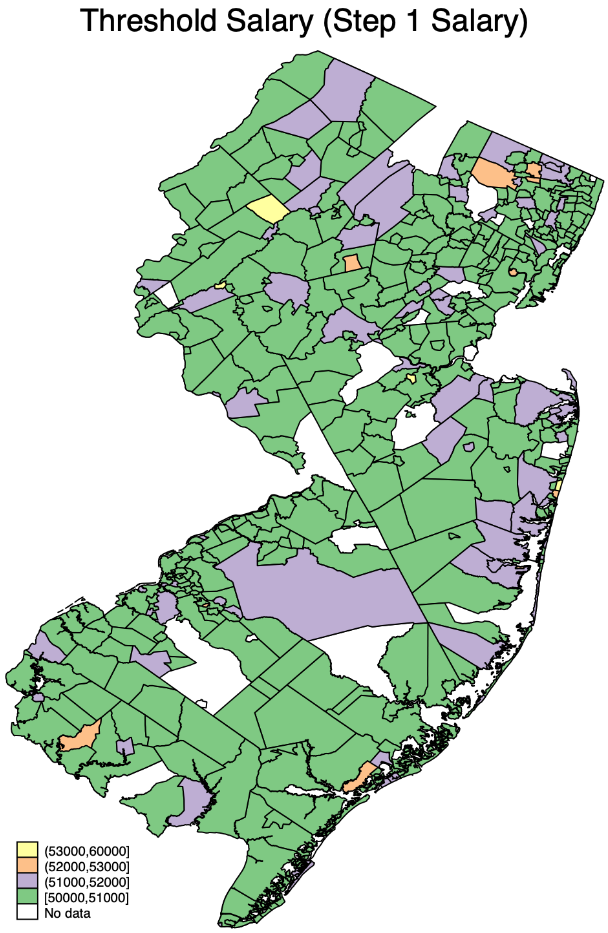

Figure 2:
Step 1 Salary

generated when a higher-paid retiring or resigning teacher is replaced by a lower-paid new hire. This allowed districts to fund salary increases without needing significant new resources. Importantly, our analysis in Section 7 finds that treated districts hired fewer staff both before and after implementing the $50K minimum salary, suggesting that districts managed wage increases partly by controlling headcount, not by increasing overall expenditures.

Although exploring the political dynamics that drove some districts to adopt the salary changes while others did not is beyond the scope of this paper, we acknowledge that broader state-level policies influenced New Jersey public schools during the study period (2003-2019). The first is the New Jersey School Funding Reform Act (SFRA), passed in 2008, which directed more state aid to districts unable to raise sufficient local revenue.[9] Another key policy was P.L. 2011 c. 78, signed by Governor Chris Christie, which increased public school staff pension contributions and significantly raised healthcare costs for teachers (New Jersey Department of the Treasury (2011)). While this policy may have reduced the appeal of teaching, we find only a small reduction in the number of teachers hired after its implementation, and

[9]In our robustness tests, we control for schools based on the socio-economic characteristics of the district they are situated in.

thus do not control for it in our main empirical specification.

3 Data

For this study, we collect data from various sources to construct three primary datasets. First, we compile data from the New Jersey Department of Education (NJDOE) website to obtain school and school district level characteristics. This includes information on enrollment, demographics, average English Language Arts (ELA) and Math scores, graduation rates, and budget expenditures. Second, we use individual teacher-level data obtained from NJDOE through a formal data request, which contains detailed characteristics for all public school teachers in New Jersey. Finally, we use data provided by the New Jersey Education Association (NJEA) to identify when each school district adopted the $50K salary policy for new teachers. Combining all of these data, we construct the following datasets.

Individual Teacher-Level Dataset

This dataset contains detailed information on all public school teachers in NJ from 2004 to 2018. It includes variables such as teacher names, job categories/codes, experience levels, and highest degrees earned. We do not have a unique identifier to track teachers across years, and thus, this dataset is cross-sectional. We primarily use this dataset to evaluate the impact of the salary negotiation campaign on teacher compensation. We also use this dataset, after employing Natural Language Processing (NLP) and Levenshtein distance, as panel data to analyze potential mechanisms (Section 10).

District-Level Dataset

This dataset spans all school districts in NJ from 2003 to 2019. It combines publicly available data from the NJDOE to assess district-level spending patterns in different areas before and after the salary negotiations.

School-Level Dataset

The ideal student outcome data for our analysis would be individual teachers' classroom scores, which we could then use to create a value-added estimate of an individual teacher's impact. However, given the unavailability of these data, we rely on public data from NJDOE.

This dataset includes data on all NJ schools from 2003 to 2019, such as student outcomes (measured as the percentage of students meeting proficiency or higher), racial demographics, enrollment, the percentage of students eligible for free/reduced lunch, total number of teachers, and other school-level characteristics. As the state examination underwent changes during this period, we use the percentage of students meeting proficiency or higher in state exams as our measure of student performance.[10] We define academic years by their Fall semesters. For example, we define the 2012–2013 year as 2012.

[10] New Jersey Department of Education. History of New Jersey's Statewide Assessments. Retrieved from https://www.nj.gov/education/assessment/history.shtml.

Our primary outcomes are Math and ELA proficiency scores for Grades 4, 8, and high school, as well as graduation rates. To ensure consistency in our analysis, we group schools into four categories: those that exclusively serve elementary (including Grade 4), those that serve middle grades (including Grade 8), those that serve high school students, and those that serve both Grades 4 and 8.

4 Empirical Strategy

We employ a staggered Differences-in-Differences (DiD) model with a continuous treatment variable to assess the impact of salary increases across New Jersey school districts. The campaign under analysis, "$50K The First Day," set minimum salary thresholds, but the exact magnitude of these thresholds varied across districts. Table 8 details the specific years when different districts negotiated the inclusion of the $50K minimum salary in their salary schedules. To capture the variation in salary levels, we treat Step 1 salaries as a continuous variable, allowing us to estimate how different levels of salary increases affect student outcomes. Our empirical strategy also leverages the variation in the timing of the negotiation approvals to identify the treatment effect.

For our primary model, we rely on the estimator proposed by De Chaisemartin and d'Haultfoeuille (2020) and De Chaisemartin and d'Haultfoeuille (2024), which is designed for settings with staggered treatment adoption. To validate the robustness of our findings, we also implement the estimators developed by Callaway and Sant'Anna (2021) and Dube et al. (2023), although these methods assume binary treatments.[11] Below, we briefly explain the structure of our empirical model, which varies slightly depending on the unit of our analysis.

4.1 Teacher-Level Model

At the teacher level, the treatment variable D_{it} represents whether the school district in which a teacher works has implemented the minimum salary threshold. We treat this as a binary indicator—set to one when the salary threshold is implemented. The dynamic DiD model for teacher-level outcomes is specified as follows:

$$Y_{it} = \sum_{t=2003}^{2019} \beta_t D_{it} + \lambda_t + \epsilon_{it} \tag{1}$$

where:

- Y_{it} is the outcome variable (e.g., teacher salary) for teacher i at time t.
- D_{it} is a binary treatment indicator that equals one in the year when the treatment is implemented.
- λ_t represents time fixed effects to account for any time-varying shocks that might affect all school districts similarly.

[11] These two estimators do not accommodate non-binary (continuous) treatments, and results using these methods are presented in the Online Appendix as robustness checks.

- ϵ_{it} is the error term.

This model allows us to estimate the impact of the new salary schedule on teacher salaries. Because teacher salaries are dependent on experience, we control (match) for experience when running this equation. We cluster errors at the school district level.

4.2 District and School-Level Model

For the district and school level, we introduce a continuous treatment variable to reflect the magnitude of salary increases across districts. In this model, D_{it} captures the intensity of the salary change (i.e., the actual Step 1 salaries in district i at time t). The dynamic model is specified as follows:

$$Y_{it} = \alpha_i + \sum_{t=2003}^{2019} \beta_t D_{it} + \lambda_t + \epsilon_{it} \tag{2}$$

where:

- Y_{it} is the outcome variable for district or school i at time t (e.g., student test scores or graduation rates).
- D_{it} is the continuous treatment variable reflecting the magnitude of the salary increase.
- α_i represents district or school fixed effects, controlling for time-invariant characteristics specific to each district/school, respectively.
- λ_t captures time fixed effects.
- ϵ_{it} is the error term.

We cluster standard errors at the school district level, as treatment occurs at this level, ensuring that our estimates account for within-district correlations over time.

Non-Binary Treatment and Estimating Event Study Estimates

In our non-binary treatment setting, we define $Y_{g,t}$ as the observed outcome for group g at time t, where $D_{g,t}$ is the continuous treatment variable indicating the Step 1 salary for school or district g at time t. The potential outcome, $Y_{g,t}(d_1, \ldots, d_t)$, represents the outcome for group g at time t, assuming the group had received treatment intensities $(d_1, \ldots, d_t)$ over time.

The treatment effect after ℓ periods of treatment, denoted $\delta_{g,\ell}$, is the difference between the actual observed outcome and the counterfactual outcome where the treatment remains at its baseline level $D_{g,1}$. Specifically, $\delta_{g,\ell}$ is defined as:

$$\delta_{g,\ell} = E[Y_{g,F_g-1+\ell} - Y_{g,F_g-1+\ell}(D_{g,1}, \ldots, D_{g,1})]$$

Here, F_g refers to the period when the school or district g first experiences a change in treatment, and $\delta_{g,\ell}$ captures the cumulative effect of treatment over ℓ periods, accounting for varying intensities of treatment.

To estimate the treatment effect, we use a Difference-in-Differences estimator $\text{DID}_{g,\ell}$, which compares the change in outcomes for treated groups to those of a control group that has not yet experienced treatment. This is expressed as:

$$\text{DID}_{g,\ell} = Y_{g,F_g-1+\ell} - Y_{g,F_g-1} - \frac{1}{N_{g,\ell}} \sum_{g'} \left(Y_{g',F_g-1+\ell} - Y_{g',F_g-1} \right)$$

In this formulation, g' indexes control groups with the same initial treatment level $D_{g,1}$ but that have not yet changed treatment by time $F_g - 1 + \ell$. The term $N_{g,\ell}$ represents the number of such control groups. This approach allows us to estimate the treatment effect for schools or districts by comparing treated groups with appropriate control groups, while taking into account the continuous nature of the treatment variable.

To complement our DiD framework, we present event study plots that visually show the dynamic effects of salary increases over time. These plots allow us to track how student outcomes evolve relative to the year when a district implements the new salary threshold. By plotting the estimated coefficients β_t over time, we can observe pre-treatment trends and test for the presence of parallel trends—an important assumption of the DiD approach. Any significant deviation from zero in the pre-treatment period would suggest that treated and control districts followed different paths prior to the policy, which could bias our estimates.

In Section 6, we present event study graphs that help visualize the impact of the $50K salary increase on student outcomes. These plots illustrate how the effects accumulate over time and how districts with different salary intensities experience varied outcomes. Readers should interpret these plots as evidence of how long it takes for the effects of salary increases to materialize and whether the improvements are sustained or diminish over time. Before showing the event study estimates, we first present the descriptive statistics.

5 Summary Statistics

We now present descriptive statistics that illustrate different school-level characteristics, student composition, and educational outcomes. These summary statistics provide a baseline understanding of the sample used in the empirical analysis, helping to identify the key features of the schools included in the study and allowing us to observe trends and differences across years.

Table 2 displays the number of schools by type. We distinguish between elementary, middle, and secondary schools, with an intermediate category for schools covering grades 4 to 8. Each school type is defined based on the grade levels it includes. Schools classified as "Four to Middle" serve students from Grade 4 through Grade 8. This category captures schools that are not traditional middle schools. This breakdown allows us to examine the differential impacts of salary increases on schools serving various age groups and grade levels. For example, high schools may face different challenges in teacher retention and performance compared to elementary or middle schools, making this classification important for interpreting the results.

Table 2: Number of Schools by Grade Type Category

School Type	Frequency
Elementary	14,076
Four to Middle	4,437
Middle	4,947
Secondary	5,185

Note: The panel data used in the main results is strongly balanced, meaning that schools are consistently observed across all years. However, school outcomes are not always measured for all of these schools.

Summary Statistics - School-Level Student Composition

Table 3 shows the composition of students at each school by year. Over the sample period, we observe a steady decline in total enrollment and student-teacher ratios. The demographic breakdown reveals that the percentage of White students has steadily declined, while the percentage of Black students has remained relatively stable.

Table 3: Detailed Summary Statistics by Year (Student Composition)

Year	Total Students	Student-Teacher Ratio	% White	% Black	Lunch Aid
2003	645.40	13.22	63.39	14.86	23.91
2004	649.44	13.13	62.45	14.96	23.89
2005	649.72	12.97	61.59	15.09	25.06
2006	648.00	12.84	60.64	15.15	24.52
2007	641.74	12.77	59.71	15.12	25.49
2008	635.82	12.61	58.88	15.04	26.16
2009	633.73	12.43	58.02	15.00	28.20
2010	631.92	12.43	57.34	14.79	29.79
2011	625.53	12.33	56.35	14.55	30.06
2012	622.32	12.77	55.58	14.35	31.58
2013	625.00	12.68	54.81	14.19	32.98
2014	621.14	12.43	54.13	14.01	34.03
2015	616.65	11.90	53.14	13.72	33.53
2016	614.48	12.02	52.19	13.52	33.38
2017	612.65	11.84	51.27	13.37	33.45
2018	609.51	11.78	50.26	13.16	32.94
2019	603.10	11.75	49.37	12.96	32.37
Total	628.60	12.46	56.42	14.34	29.51

Note: The "Lunch Aid" column represents the percentage of students eligible for free or reduced-price lunch. The fourth and fifth columns denote the racial composition of the student body, focusing on the percentage of White and Black students, respectively.

Summary Statistics - School-Level Educational Outcomes

Tables 4 and 5 summarize key educational outcomes used in this study. The data reveals significant variation across years, particularly after 2015, when New Jersey's core standards were revised. This revision aimed to set consistently high expectations for student performance, leading to noticeable changes in standardized test outcomes post-2015. This drastic drop in proficiency scores could be problematic, and thus, we also run our estimates restricting our analysis to years before such drastic drops are observed. Our analysis remains robust when restricting the sample to pre-2015 data and using standardized scores by year, ensuring that the results are not driven by changes in educational standards alone.

Table 4: Summary Statistics by Year (Educational Outcomes)

Year	**Math 4**	**Math 8**	**Math HS**	**ELA 4**	**ELA 8**	**ELA HS**	**Grad.**	**Post Enroll.**
2003	70.44	60.04	67.43	82.58	75.52	81.71	86.02	82.58
2004	74.21	64.58	71.47	85.79	73.08	83.72	85.94	83.46
2005	81.38	64.71	76.58	84.79	73.13	84.33	87.03	84.19
2006	83.42	66.41	77.05	83.34	75.01	84.86	89.09	84.40
2007	85.66	70.66	74.36	83.90	74.12	86.37	89.82	86.06
2008	85.63	68.99	75.86	85.37	82.04	84.11	89.58	85.58
2009	74.44	71.95	73.63	68.01	82.32	84.62	89.63	84.79
2010	78.43	68.99	74.96	64.38	82.61	88.18	90.70	86.09
2011	80.60	71.97	76.14	67.88	82.26	90.54	88.60	.
2012	79.14	73.00	83.03	64.43	82.61	91.72	88.65	.
2013	79.99	70.38	84.48	65.46	81.90	92.20	89.35	76.19
2014	76.53	71.76	83.64	65.64	79.67	91.79	89.88	76.68
2015	44.35	30.04	42.11	57.39	53.85	38.50	90.71	77.11
2016	49.87	31.55	44.52	59.46	56.28	44.92	91.22	77.73
2017	50.42	32.77	47.37	61.88	59.48	47.15	91.29	76.64
2018	52.08	33.58	47.18	63.44	60.89	51.38	91.73	78.69
2019	53.69	35.12	50.39	62.72	63.78	58.24	91.86	77.58
Total	70.77	59.53	67.59	70.96	72.98	75.79	89.48	81.18

Note: "Grad." denotes high school graduation rates. "Post Enroll." refers to the percentage of students enrolling in post-secondary education within 16 months of graduation. In 2015, New Jersey's core standards were revised, impacting math and ELA test scores. Our results remain robust when restricting the analysis to data from years prior to 2015.

Table 5: Summary Statistics for Educational Outcomes

Variable	**Observations**	**Mean**	**Std. Dev.**	**Min**	**Max**
Math 4	17,945	70.77	21.69	5.6	100
Math 8	8,187	59.53	24.44	3.5	100
Math HS	5,203	67.59	26.01	7.4	100
ELA 4	13,848	70.96	18.79	10	100
ELA 8	9,587	72.98	20.17	10	100
ELA HS	5,134	75.79	22.72	10	100
HS Graduation	5,189	89.48	12.12	4.6	100
College Graduation	4,589	81.18	12.02	16.6	100

Given the detailed overview of the school characteristics, student composition, and educational outcomes in this section, we now begin discussing our results.

6 Effects of the Campaign on Salaries

In this section, we show the effects of the campaign on staff salaries. We begin by estimating Equation 1, restricting our sample to only Math and ELA teachers. Our treatment here is binary.

Figure 3: Salary of Math and ELA Teachers

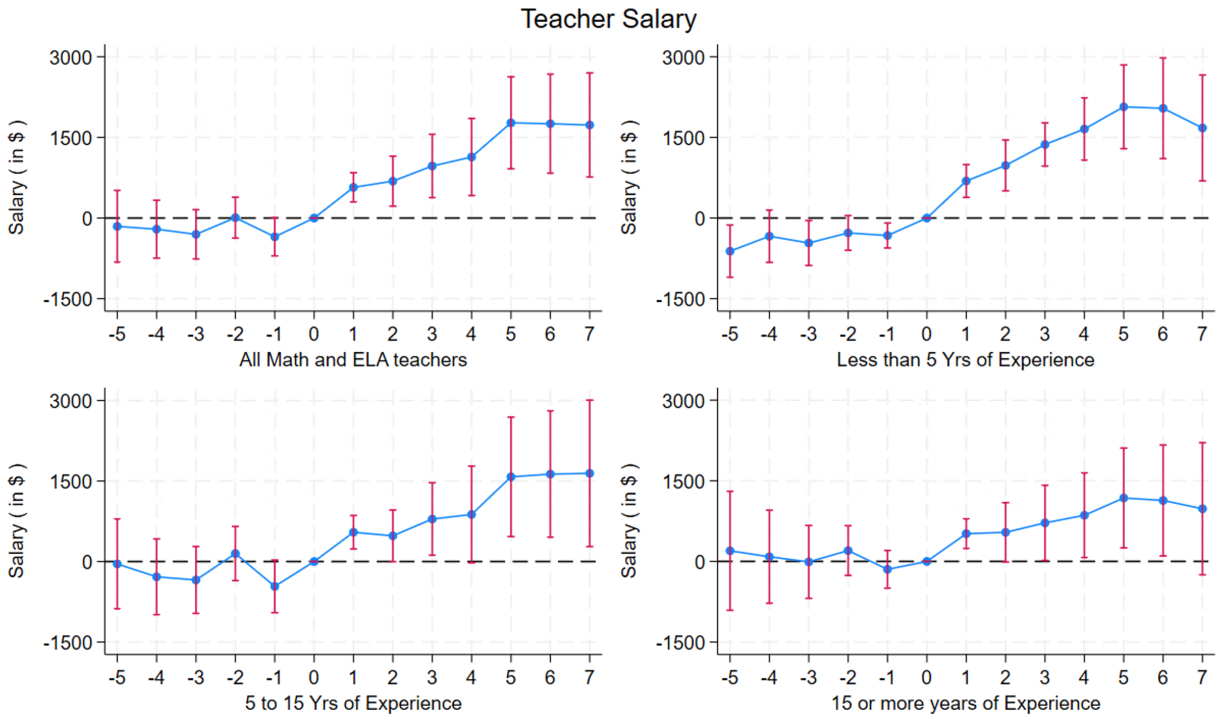

Note: Point estimates and 95% confidence intervals of parameters $\beta_t D_{it}$ in Equation (1). The event study shows how teacher salaries changed over time relative to the year the treatment (the salary campaign) was introduced. The x-axis represents the number of years before and after the salary increase took effect (with "0" being the year the policy was implemented), while the y-axis shows the estimated change in salaries. Each blue line represents the difference in salary for treated teachers compared to not-yet-treated teachers in that given year. The red bars show the 95% confidence intervals.

The first graph in Figure 3 (top left) presents an event study of all Math and ELA teachers. This event study shows how teacher salaries changed over time relative to the year the treatment (the salary campaign) was introduced, aggregating across all teachers regardless of experience level. The x-axis represents the number of years before and after the salary increase took effect (with "0" being the year the policy was implemented), while the y-axis shows the estimated change in salaries. Each blue line represents the difference in salary for treated teachers compared to not-yet-treated teachers in that given year. The red bars show the 95% confidence intervals.

Salaries begin to increase immediately after treatment (year 0) and continue to rise consistently in the following years. Before treatment (to the left of 0), we see that the differences in salaries between treated and control groups were nearly identical—this supports the idea that treated and not-yet-treated schools were on parallel trends prior to the policy. We then break down our event study analysis based on the level of experience.

We now examine teachers with less than five years of experience (top right panel in Figure 3). These new teachers in treated schools see a significant salary increase relative to their peers in not-yet-treated schools. However, the parallel trends assumption could appear to be slightly violated for this group, as pre-treatment salaries show a small upward movement in treated schools before the salary floor was officially implemented. This could raise concerns about the validity of the results for this subgroup. To address these concerns, we apply the methods illustrated by Rambachan and Roth (2020) for robust inference in difference-in-differences and event study designs when the parallel trends assumption may be violated.[12]

When we break down the results by other experience groups (top right and bottom left panels), we observe positive salary increases for teachers with 5 to 15 years of experience and teachers with over 15 years of experience. For these groups, the pre-treatment trends align closely between treated and control schools, suggesting that the parallel trends assumption is satisfied. These teachers see salary increases averaging around \$1,000 to \$2,000, confirming that the salary campaign benefits teachers at all stages of their careers. This suggests that the policy successfully improved teacher salaries without creating potential discontent among senior staff members, who might have otherwise felt disadvantaged by a campaign that only benefited new hires. While breaking down the results by experience level (as seen in the other panels) provides valuable insights, the all-teachers analysis is particularly important because it allows us to assess the overall impact of salary increases across the entire group of teachers. This broader view helps us examine whether increasing salaries for all teachers could potentially improve student outcomes.

We next run the same analysis excluding Math and ELA teachers and including all other staff members (both teaching and non-teaching). The top panel in Figure 4 shows that other teaching staff members, such as science, history, or arts teachers, also experience similar positive impacts on their salaries after the campaign. We then restrict the sample to include only non-teaching staff (e.g., administrators, support staff). The event study (bottom panel) shows that these groups also experienced no negative changes in their salaries. Salary differences remained constant throughout the period of analysis. Taking stock, the event study results show that salaries increased for all subgroups, with new teachers experiencing salary increases the very year the negotiation was approved. Teachers with more experience saw increases the year after the act, and overall, no teachers were disadvantaged. Table 9 in the Appendix shows the average impact of the \$50K The First Day Campaign on salaries for all full-time staff.

[12] Using this approach, we find that even when accounting for potential deviations in pre-treatment trends, the salary increases for new teachers are significantly higher than would be expected under linear extrapolation of pre-treatment trajectories. This provides strong evidence that the salary increases are indeed driven by the campaign and not by underlying trends.

Figure 4: Salary of Teaching Staff Outside Math/ELA

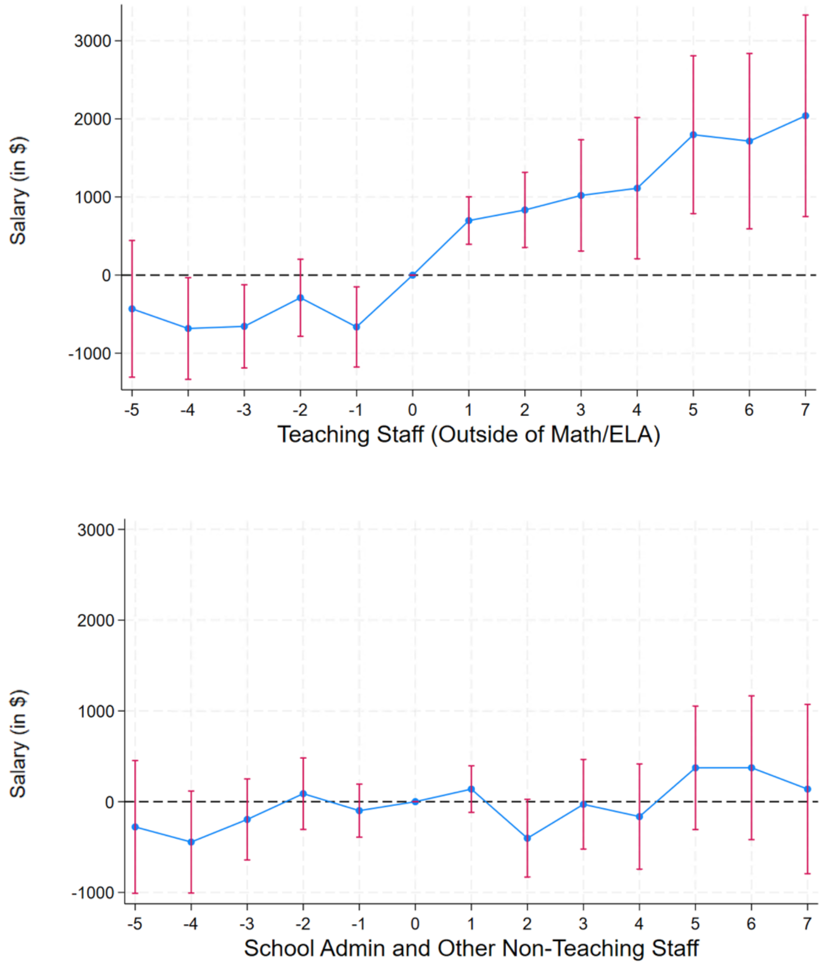

Note: Point estimates and 95% confidence intervals of parameters $\beta_t D_{it}$ in Equation (1). The y-axis shows yearly salary, and the x-axis refers to the time period before and after the treatment. The reference year is time period 0, which indicates the year the treatment took place.

Taken together, these findings allow us to rule out any negative spillover effects among different types of staff due to the salary increases for teachers. The absence of salary reductions or stagnation for other staff members, both teaching and non-teaching, suggests that the policy was implemented in a way that avoided internal disparities within the district's workforce. This minimizes the risk of unintended consequences, such as dissatisfaction among non-teaching staff or other educators not directly targeted by the salary floor.

7 Effects on District Spending and School Composition: Where is the Money Coming From?

In this section, we examine the impact of the salary campaign on district-level spending and compositional changes at the school level. We begin by estimating Equation 2 at the district level to analyze spending patterns within a district before and after the campaign.

Figure 5: School District Spending (in Dollars)

Note: Point estimates and 95% confidence intervals of parameters $\beta_t D_{it}$ in Equation (2). Unit of analysis is at the school district level. The y-axis shows yearly spending, and the x-axis refers to the time period before and after the treatment. The reference year is time period 0, which indicates the year the treatment took place. Treatment is non-binary (continuous). The figure on the top left shows per-pupil spending in all areas. The figure on the top right shows per-pupil spending used only on teacher salaries. The bottom figure shows per-pupil spending outside of expenditures on classroom salaries.

Figure 5 shows the event study for per-pupil spending at the district level. The top-left panel demonstrates that total per-pupil spending did not significantly increase in treated schools, suggesting that districts did not receive or allocate additional overall budget resources as a result of the treatment. The bottom panel further indicates that differences in spending on non-salary components remained largely unchanged, implying that changes in spending outside of teacher salaries could not have driven the observed outcomes. The top-right panel suggests that differences in per-pupil spending on teacher salaries did not increase post-treatment. This may initially appear counterintuitive given that teacher salaries increased. However, this result can be explained by several factors.

Firstly, the difference in the total amount spent per student on teacher salaries between treated and not-yet-treated schools did not change significantly due to shifts in staffing patterns. Specifically, the number of staff hired post-treatment decreased in treated districts, as shown in Figure 6. This compositional change, where schools hired fewer teachers or replaced departing, higher-paid senior staff with lower-paid new hires, helps explain why total spending per student on salaries remained relatively flat, even though individual teacher salaries increased. This phenomenon arises from the fact that while individual salaries rose, the number of teachers in treated districts decreased, balancing out the total spending per student.

Figure 6: School-Level Changes

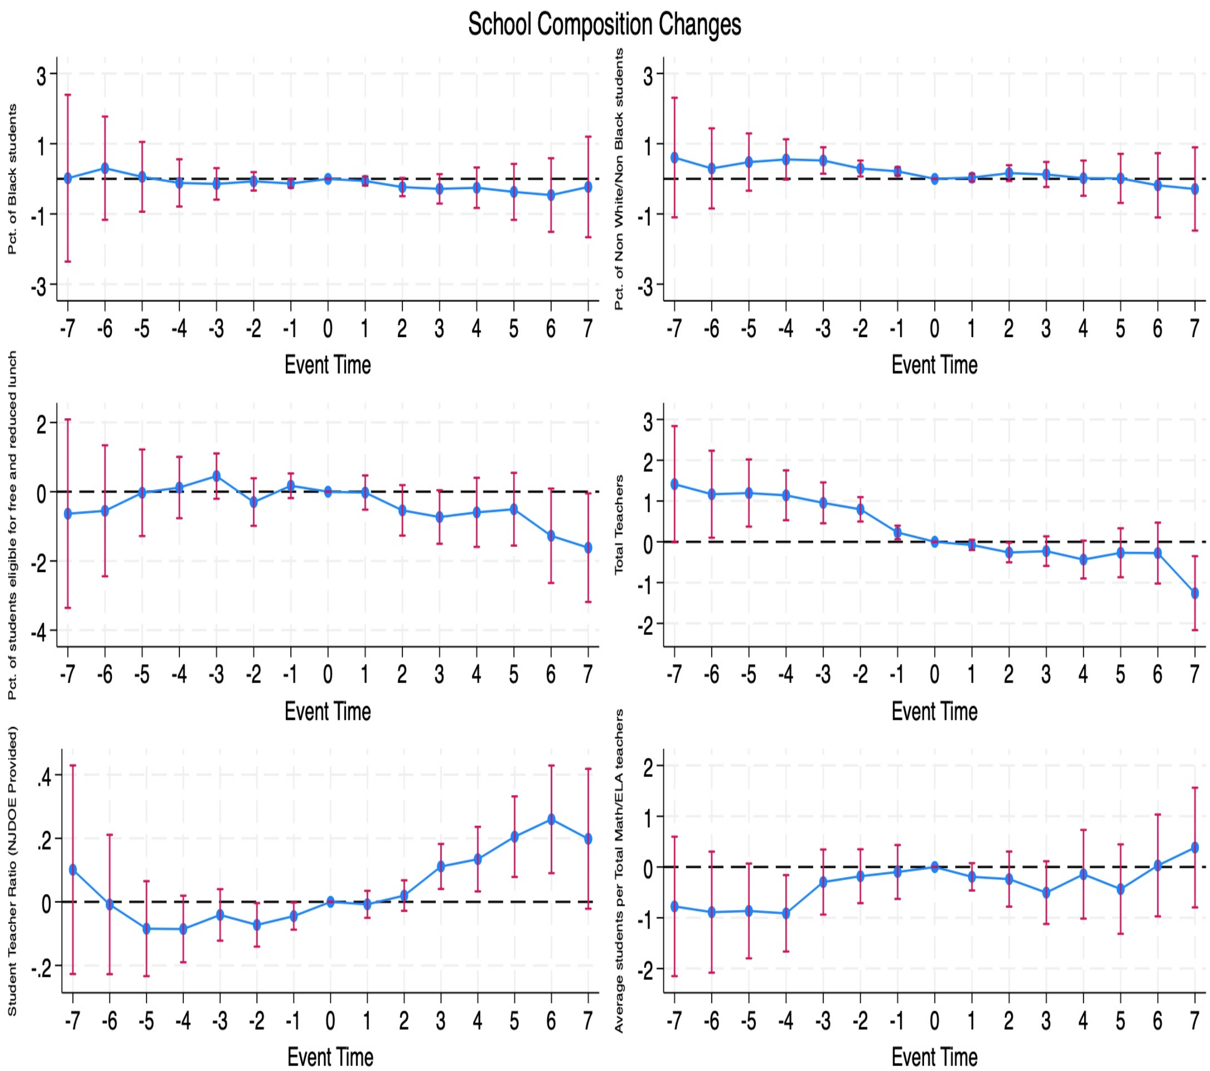

Note: Point estimates and 95% confidence intervals of parameters $\beta_t D_{it}$ in Equation (2). Unit of analysis is at the school level. The y-axis shows yearly outcomes, and the x-axis refers to the time period before and after the treatment. The reference year is time period 0, which indicates the year the treatment took place. Treatment is non-binary. The racial composition of students remains fairly similar, and the total number of staff members employed (bottom right) decreases.

Figure 6 also highlights the compositional changes in student demographics at the school level. The racial composition of students (top panels) remains consistent before and after the campaign, suggesting no significant demographic shifts that could confound our analysis. The percentage of students eligible for free and reduced lunch remains fairly stable, with only about a 1% drop in treated schools. Before treatment, treated schools had more teachers than not-yet-treated schools, which explains why total spending per student on salaries remained relatively flat.

Hiring fewer teachers could increase the burden on teachers in treated schools, and we do observe a slight increase in the student-teacher ratio in treated schools post-treatment (bottom left panel). Although this increase is modest (less than one additional student per teacher), an increase in the student-teacher ratio could potentially affect student outcomes. If this compositional change were to influence student outcomes, we argue that it would likely introduce a downward bias. Following the findings from Angrist and Lavy (1999), a higher student-teacher ratio generally leads to lower student outcomes, making any potential improvements in student performance even more noteworthy given this slight increase. Importantly, the ratio of students to Math and ELA teachers—the subjects directly targeted by the salary increases—remains stable throughout the study period. This suggests that Math and ELA teachers in treated schools were not overburdened with more students after the campaign.[13]

Overall, the results presented here address several key concerns about resource reallocation and staffing composition changes that could potentially confound our later analysis of student outcomes. Specifically, while individual teacher salaries increased, the total per-pupil spending on salaries remained stable due to reductions in staffing levels, effectively mitigating the risk of budgetary distortions. Additionally, the stability in the ratio of students to Math and ELA teachers suggests that core subject teachers were not overburdened, minimizing the risk of increased workloads that could negatively impact student performance.

Moreover, potential endogeneity issues—such as schools adjusting their spending patterns in other areas or experiencing demographic shifts—are largely ruled out. The compositional stability of the student body and the lack of significant changes in district spending on non-teaching resources ensure that our analysis remains focused on the causal effect of salary increases rather than confounding factors. These findings lend credibility to the robustness of our identification strategy, allowing us to more confidently assess the impact of teacher salary increases on student outcomes. With these considerations in place, we now move on to examine whether higher salaries lead to improvements in student performance.

8 Main Results - Impact of the Campaign on Student Outcomes

This section presents the main results of the paper where we estimate Equation 2 using individual schools as our unit of analysis. The key outcomes of interest are Math and ELA proficiency levels across grades 4, 8, and high school, as well as graduation rates and college enrollment.

[13] Although not shown for brevity, we conducted the event study on other areas of spending as well and found no meaningful changes due to this campaign.

8.1 Math Outcomes

We begin by examining the impact of the salary increases on Math proficiency outcomes, with the results illustrated in Figure 7. This figure presents an event study of Math proficiency levels for 4th grade, 8th grade, and high school, providing a detailed look at how these outcomes evolved before and after the salary increases were implemented.

Figure 7: Proficiency Level Met (Math)

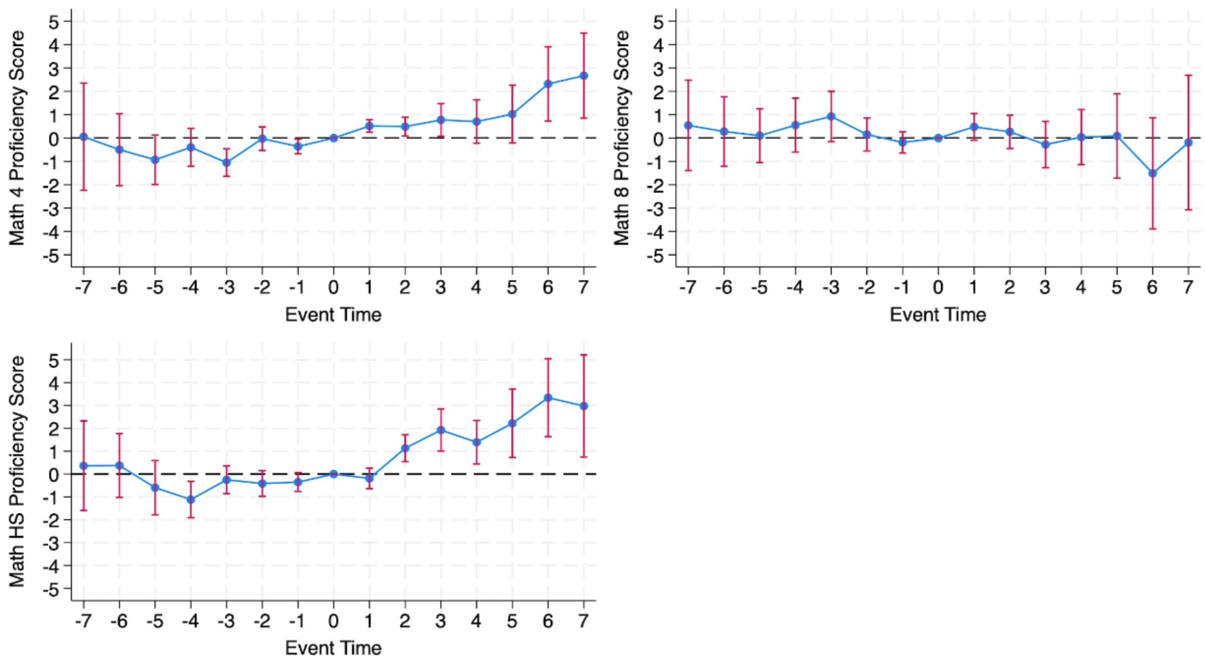

Note: Point estimates and 95% confidence intervals of parameters $\beta_t D_{it}$ in Equation 2 at the school level. The y-axis shows changes in proficiency level, and the x-axis refers to the time period before and after the treatment. The reference year is time period 0, which indicates the year the salary campaign took place. Treatment is non-binary.

In interpreting the event study estimates, the horizontal axis represents time relative to the introduction of the salary increases (with year "0" being the treatment year), and the vertical axis represents changes in the percentage of students meeting Math proficiency standards. The blue dots in the figure show the estimated treatment effect at each time period, while the red bars represent the 95% confidence intervals around these estimates. A value of 0 on the y-axis would indicate no change in differences in proficiency levels, while positive values suggest improvements in Math performance following the salary increases.

The pre-treatment period (to the left of year 0) shows that the estimated treatment effects are close to zero, which validates the parallel trends assumption—i.e., the proficiency levels in treated and control schools were evolving similarly before the policy was implemented. This is crucial for establishing that the post-treatment effects can be attributed to the salary increases rather than to pre-existing differences between treated and not-yet-treated schools.

After the salary campaign (year 0 and beyond), we observe a significant and sustained increase in Math proficiency, particularly in 4th grade and high school. The pattern for high school Math is consistent with that of 4th grade, demonstrating that the salary increases benefited students across multiple educational levels. Importantly, the effects appear to persist over time, with continued improvements in proficiency levels even several years after the salary increases were implemented. We do not observe a significant increase in 8th-grade scores post-treatment, and the estimates for 8th grade exhibit larger standard errors. These findings highlight the differential effects of the salary campaign across different grade levels, with the most pronounced improvements observed in early and late stages of schooling.

8.2 ELA Outcomes

We next turn to the event study results for ELA outcomes, as shown in Figure 8.

Figure 8: Proficiency Level Met (ELA)

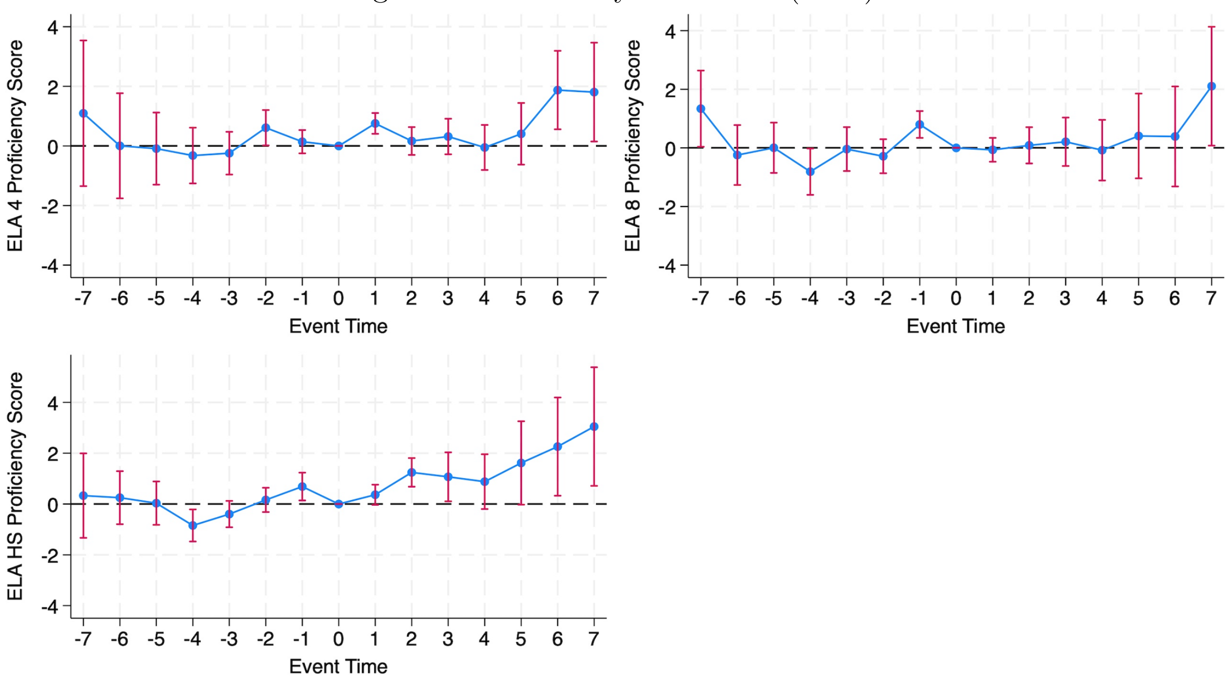

Note: Point estimates and 95% confidence intervals of parameters $\beta_t D_{it}$ in Equation 2 at the school level. The y-axis shows changes in proficiency level, and the x-axis refers to the time period before and after the treatment. The reference year is time period 0, which indicates the year the treatment took place. Treatment is non-binary.

The vertical axis shows changes in the percentage of students meeting ELA proficiency standards, where positive values reflect improvements in performance. The pre-treatment period (left of year 0) shows that ELA proficiency levels were evolving similarly in treated and not-yet-treated schools before the salary increases were introduced.

Following the implementation of the salary increases, the event study reveals a significant and sustained positive effect on ELA proficiency in both 4th grade and high school. The timing of the improvement in high school ELA is similar to that of 4th grade, with proficiency rates rising immediately after the salary campaign and continuing to increase over subsequent years. These results suggest that the policy's benefits extended across different educational levels, with students at both ends of the schooling spectrum benefiting from the enhanced teacher compensation. However, as with the Math outcomes, the event study for 8th-grade ELA shows no significant impact. The estimates for 8th-grade ELA proficiency exhibit larger standard errors, making it difficult to detect a clear pattern.

8.3 Graduation and College Enrollment

Figure 9 examines the effects of salary increases on high school graduation rates and college enrollment. This analysis allows us to assess the potential long-term educational attainment impacts beyond proficiency scores.

Figure 9: Graduation and College Enrollment

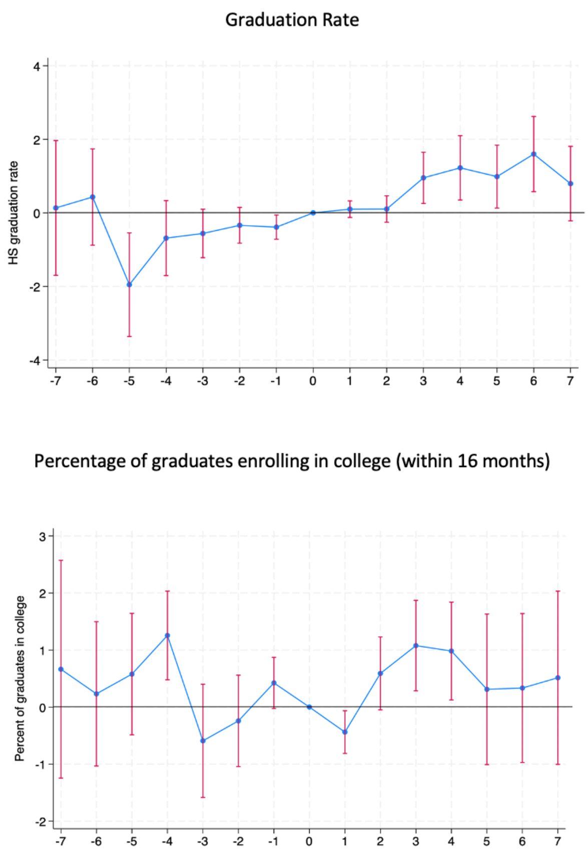

Note: Point estimates and 95% confidence intervals of parameters $\beta_t D_{it}$ in Equation 2 at the school level. The y-axis shows changes in graduation rates and college enrollment, respectively. The x-axis refers to the time period before and after the treatment. The reference year is time period 0, which indicates the year the treatment took place. Treatment is non-binary.

From the event study in Figure 9, we observe a positive and significant impact of the salary increases on high school graduation rates. However, we do not see a similar rising trend in the percentage of graduates enrolling in college. Although the salary campaign may have created a more conducive learning environment, the long-term effects on college enrollment are not as pronounced.

8.4 Taking Stock

In this section, we summarize the key findings by quantifying the overall impact of the salary increases across student outcomes. Table 6 provides the average treatment effects over the seven-year period following the implementation of the salary campaign.

Table 6: Estimation of Treatment Effects: Average Total Effect

Score	Estimate	SE	N	Switchers	Estimate/SD	Estimate/Mean
Math 4	1.17	0.45	10,305	4,231	0.05	0.017
Math 8	-0.0785	0.625	6,000	2,293	< 0.01	< 0.01
HS Math	1.72	0.49	3,606	1,459	0.06	0.026
ELA 4	0.73	0.39	10,339	4,255	0.03	0.010
ELA 8	0.37	0.51	6,911	2,746	0.02	0.001
ELA HS	1.48	0.56	3,767	1,596	0.06	0.010
Grad. Rate	0.93	0.31	3,807	1,629	0.07	0.005
College Grad. Rate	0.44	0.41	2,850	982	0.04	0.001

Note: Table reports the average treatment effect of implementing the new salary guide on student outcomes. N represents the total number of observations (unbalanced). The fourth column reports the total number of schools that switched treatment across the time period. The fifth column shows the ratio of estimate to standard deviation, and the last column shows the ratio of estimate to mean scores.

The estimates are presented in terms of standard deviations (SD) and average scores. We briefly discuss the magnitude of the effects by analyzing the estimates.

For 4th-grade Math, scores increase by approximately 0.055 SD on average, reflecting a substantial improvement in the proportion of students meeting proficiency standards. High school Math proficiency shows a similar positive trend, with scores increasing by about 0.06 SD post-treatment. In terms of the ratio to average scores, math scores for 4th grade and high school increased by roughly 0.02.

The results for ELA proficiency are also positive, though slightly smaller in magnitude compared to Math. In 4th-grade ELA, the average increase is around 0.045 SD, while high school ELA proficiency improves by approximately 0.06 SD. In terms of the ratio to average scores, ELA scores for 4th grade and high school increased by roughly 0.01.

Regarding graduation rates, we observe a meaningful increase of 0.06 SD post-treatment. This suggests that the salary increases not only impacted student performance in standardized tests but also had a lasting effect on high school completion, a critical indicator of long-term educational success. The impact on college enrollment, while positive, is less conclusive. We estimate an effect of around 0.04 SD, but the estimates become noisier in the years following the salary increases. In terms of the ratio to average graduation rates, the increase was relatively modest at 0.003.

Overall, the evidence suggests that increasing teacher salaries led to meaningful improvements in student outcomes, with the most pronounced effects seen in 4th grade and high school, while 8th-grade outcomes remained largely unchanged. The magnitude of these effects suggests that improving teacher salaries could be a potentially powerful tool for boosting educational performance across different stages of schooling.

8.5 Robustness Checks, Alternate Specifications, and Heterogeneous Effects

To ensure the robustness of our findings, we apply several alternative specifications to verify the consistency of our results.

First, we restrict the sample to a strongly balanced panel, focusing only on schools that have complete data for all event-time periods. In line with the DiD literature, we refer to this approach as restricting to "same switchers." By including only those schools with data for all specified event-time periods, we avoid the potential for compositional bias that could arise if different schools contribute to different periods. This approach ensures that the treatment effects are estimated based on a consistent set of schools over time, which strengthens the internal validity of our results. However, this restriction reduces the sample size and could introduce selection bias, as the final sample may no longer be representative of the broader population of schools. Figure 12 in Appendix B shows that our results remain consistent under this more restrictive specification.

Next, we standardize the outcome variables to account for changes in testing standards that were introduced in 2015. This standardization ensures that any observed effects are not confounded by shifts in assessment criteria, allowing for a more accurate comparison of student performance over time. Figure 13 in Appendix B illustrates the event study estimates using standardized outcome values.

Lastly, we employ two alternative methods for estimating the event study effects. The first, proposed by Callaway and Sant'Anna (2021), uses a group-time average treatment effect approach under a binary treatment setting. The second, introduced by Dube et al. (2023), uses local projections to estimate the treatment effects, providing a different way to handle staggered adoption by projecting the effects over time. Appendix C shows the event study of the outcomes under these approaches. Our results remain robust across all specifications.

To explore potential heterogeneity in the effects, we estimate the model separately for schools grouped by the socio-economic status (SES) of their districts. This approach enables us to assess whether the impact of salary increases differs across schools with varying SES backgrounds. However, the estimates obtained from this disaggregated analysis are too noisy to draw definitive conclusions about differential effects based on socio-economic status since schools are divided into seven distinct groups.

We do control for SES in our robustness test, by using broader groupings of schools according to their relative socio-economic status, categorized based on the district factor group of the school districts. Under this broader classification, where schools are matched with other schools in the same socio-economic group, the results remain consistent, and in many cases, the estimates are more precise (see Appendix D).

We also run our estimates by randomizing the year treatment took place by entering fake years of treatment. Appendix E shows the event study for student outcomes where the treatment year is set to be three years later than the actual treatment year, and the treatment intensity is randomized.

9 Comparison with Other Studies

In this section, we compare the magnitude of our estimates with findings from other relevant literature. Our study focuses on a constant increase in salaries for all public school teachers in New Jersey. It is essential for readers to note that this approach differs from studies that evaluate performance-based salary schedules. Since we do not distinguish between teachers, salary increases occur for all teachers, irrespective of their performance.

In two 2013 studies, Fryer (2013) and Goodman and Turner (2013) found that bonuses of $1,500 to $3,000 per teacher in New York City public schools had little impact on teacher effort, student performance in math and English, or classroom activities. In another study from Tennessee, Springer et al. (2012) found that students whose teachers were eligible to receive bonus payments performed at the same level as those whose teachers were ineligible, indicating no significant effect. Similarly, Biasi et al. (2021) noted that after Act 10 in Wisconsin, wage growth remained small and negative for teachers who stayed in their positions, but increased significantly for those who moved to new districts ($1,750 at the median). This led to flexible pay (FP) districts increasing reading and math scores by 0.4 and 0.6 standard deviations, respectively. However, salary changes varied in this study, and thus a direct comparison with our results may not be appropriate.

Baron (2018) examined the impact of Act 10 on average student achievement in Wisconsin, finding that the reduction in union power decreased teacher salaries by roughly 4% and reduced average Wisconsin Knowledge and Concepts Examination (WKCE) scores by approximately 20% of a standard deviation. We begin by comparing the magnitude of our estimates with these two Wisconsin-focused studies. Before delving into this comparison, we present the summary statistics of teacher salaries by year in Table 7.

Table 7: Summary of Salary by Year with Median, SD, and Number of Observations

Year	**Mean Salary**	**Median Salary**	**Std. Dev.**	**N**
2004	60,274	54,478	19,198	124,898
2005	61,882	55,655	19,471	126,748
2006	63,571	57,248	19,684	123,361
2007	65,264	58,936	19,705	125,249
2008	67,082	60,963	19,760	125,978
2009	68,992	63,239	19,735	126,319
2010	70,422	65,211	19,569	120,526
2011	70,823	65,643	19,383	120,725
2012	71,186	65,851	19,580	120,594
2013	72,077	66,998	19,596	123,871
2014	72,456	67,451	19,527	124,597
2015	73,051	68,200	19,503	123,485
2016	74,237	69,743	19,544	124,902
2017	75,436	71,375	19,592	123,971
2018	76,397	72,520	19,612	125,159
Total	69,523	65,188	20,141	1,860,383

Note: Table reports the average and median salary for all full-time staff in NJ public schools rounded to integer values. Individual teacher-level salary data are winsorized at the 1st and 99th percentiles. The fifth column reports the total number of full-time staff (both teaching and non-teaching) in NJ public schools.

From Table 9, which shows the event study estimates of the campaign on all staff in a school, we observe that differences in salaries between teachers in treated and control schools increased by approximately $1,200 on average. This corresponds to roughly 1.7% of the average staff salary and around 1.9% of the median staff salary. Taking an average of all scores in Table 6, excluding graduation rates and 8th-grade scores, we find that a 2% increase in salaries led to an approximate increase in proficiency levels of 0.05 SD (5% of a standard deviation). This is very similar to the effects observed in Biasi (2021) and slightly lower compared to the effects observed in Baron (2018).[14]

Given that not all districts may be able to achieve these salary increases through staff restructuring alone, we now estimate the additional budget required to achieve a uniform $2,000 salary increase for all staff. This amount of salary increase is double the average salary increase for all staff but equal to the highest estimates for teaching staff (Figures 3 and 4). According to current data from the NJDOE New Jersey Department of Education (2024), there are approximately 600 students per school and 50 full-time teachers. At $2,000 per teacher, this would require an additional $100,000 per school, translating to an increase in per-pupil spending of roughly $170. If this amount were to be collected through increased government spending, assuming 1.4 million students attend public schools, the additional

[14] It is important to note that the effect of a decrease in teacher salaries on student achievement may not necessarily mirror that of an increase in teacher salaries.

budget needed would amount to approximately \$250 million. This represents roughly 2.4% of state aid for education and about 0.03% of New Jersey's Gross Domestic Product (GDP).

To put this in context, Jackson and Mackevicius (2024) finds that increasing school spending by \$1,000 per pupil (sustained over four years) raises test scores by 0.03 SD and increases college enrollment by 2.8 percentage points. Based on our results, increasing per-pupil spending on teacher salaries by just \$170 yields similar improvements in test scores (around 0.03 SD), suggesting a higher short-term return on investment. However, this comparison should be interpreted cautiously, as the relationship between spending and outcomes may not be strictly linear, especially at higher levels of spending. Nonetheless, under a linear assumption, increasing per-pupil spending on teacher salaries by \$1,000 could potentially lead to a proficiency score increase of 0.15 SD and raise college enrollment rates by around 4 percentage points. We now explore the mechanisms through which these positive effects could have been observed.

10 Exploring Potential Mechanisms: Are Teachers Switching?

Higher salaries could improve teacher performance through different mechanisms. First, increased pay may enhance productivity by raising the stakes of potential dismissal, which is in line with the efficiency wages model. Alternatively, higher salaries could improve job satisfaction, thereby motivating better performance. The second mechanism suggests that offering higher salaries enables districts to attract higher-quality teachers, either by encouraging effective teachers to switch schools or by drawing more talented individuals into the teacher labor market.

From a policy perspective, an ideal outcome would avoid districts competing for teachers, as this could exacerbate achievement gaps between schools. Moreover, a district can only afford to dismiss underperforming teachers if it has a robust pool of candidates to replace them. We now turn to an examination of what might be driving our positive results, weighing the evidence for and against each of these two mechanisms.

Our analysis is inspired by findings from Biasi et al. (2021) and Biasi (2021). In these two papers, the authors show that granting districts control over teacher pay leads to more efficient but also more unequal teacher distribution. Efficiency improves as districts can better reward teachers for their contributions, encouraging sorting based on comparative advantage. However, inequality worsens, as teachers tend to prefer working in districts with high-achieving students. Flexible pay policies make it easier for wealthier districts to attract these teachers, thereby widening the gap. We aim to see if our data reflect a similar pattern. We provide a brief theoretical overview of this possibility in the Appendix (Section A.1).

Approach Used for Mechanism Analysis

Ideally, our teacher dataset would include a unique identifier for each teacher to track their movement over time. However, our data lack such a variable.[15] Instead, we rely on first name, last name, and date of birth (DOB) to track teachers. This method poses challenges due to name changes (e.g., from marriage) and frequent spelling or DOB errors. To address these

[15] We have submitted a request to the NJDOE for data containing unique teacher identifiers.

issues, we employ Levenshtein distance and Natural Language Processing (NLP) algorithms to match teacher names. We incorporate NLP because using only Levenshtein distance would fail to match common nicknames (e.g., Robert to Bob).[16]

We briefly describe the matching process. First, after restricting our sample to only Math and ELA teachers, we match combined name and DOB to create unique IDs using Levenshtein distance. However, due to the absence of DOB data after 2017, our analysis based on DOB is restricted to the years 2004–2016. For teachers with fewer than five years of records, we then match using only first names to account for potential changes in last names and missing DOBs. We rely on NLP for this approach. On average, we generate seven years of panel data per teacher, which we argue is sufficient to detect teacher movements and new hires. We classify teachers into two categories: "switchers" and "new teachers." Switchers are defined as those who moved to a new district after working in another district the previous year, while new teachers have fewer than three years of experience and did not previously switch districts. Though this method is not foolproof, the seven-year panel should capture relevant shifts. Consequently, our event study focuses on a narrower window to minimize the influence of fewer observations at the margins, where confidence intervals would be wider due to limited data.

Patterns in Elementary Schools

We start by analyzing teacher movement patterns for Math and ELA teachers in elementary schools and find fewer than 15 occurrences of switching throughout the time period. Thus, for elementary school teachers, we do not run an event study on the number of switchers before and after the campaign. Figure 10 shows the results for elementary teachers.

[16] Our NLP approach matched fewer than 15 individuals.

Figure 10: Mechanism Elementary

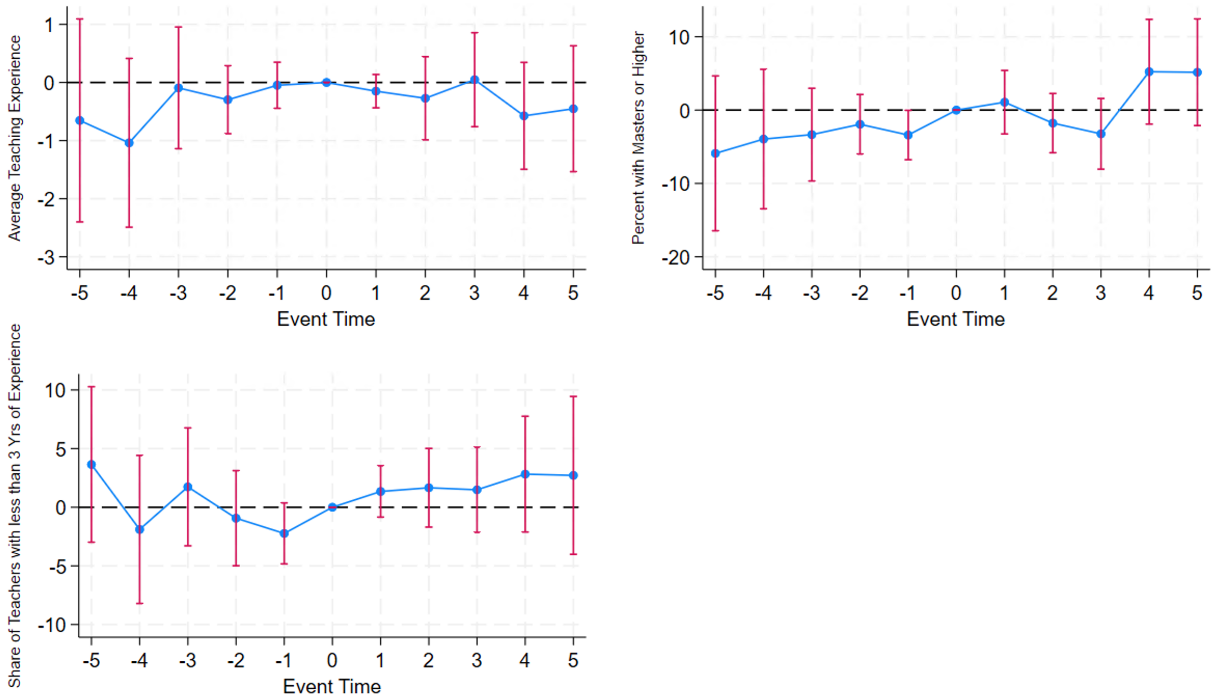

Note: Point estimates and 95% confidence intervals of parameters $\beta_t D_{it}$ in Equation (2). Average experience, total teachers with a Master's Degree or higher, and total new teachers (< 3 years of experience) calculated from aggregating teacher-level data up to the school level. Treatment is continuous. The teacher dataset used incorporates NLP and Levenshtein distance.

From Figure 10, we can rule out any significant teacher compositional effects. Under the assumption that our teacher matching algorithm was unsuccessful, if teachers did switch post-policy, we would expect these schools to have a higher average teaching experience. However, we fail to observe this from the first figure. The top right figure then shows that the schools also did not hire more teachers with a Master's degree or higher. The bottom figure runs the event study using the share of new teachers as the outcome. Here too, we find no consistent difference in hiring patterns before or after the campaign. In summary, from these sets of event studies, we rule out that treated schools recruited "higher quality" teachers from other schools.

Patterns in High Schools

We now analyze teacher movement patterns for Math and ELA teachers in high schools. We observe 1,620 occurrences of switching, and thus include the share of switchers in our analysis. The event study estimates are shown in Figure 11.

Figure 11: Mechanism High School

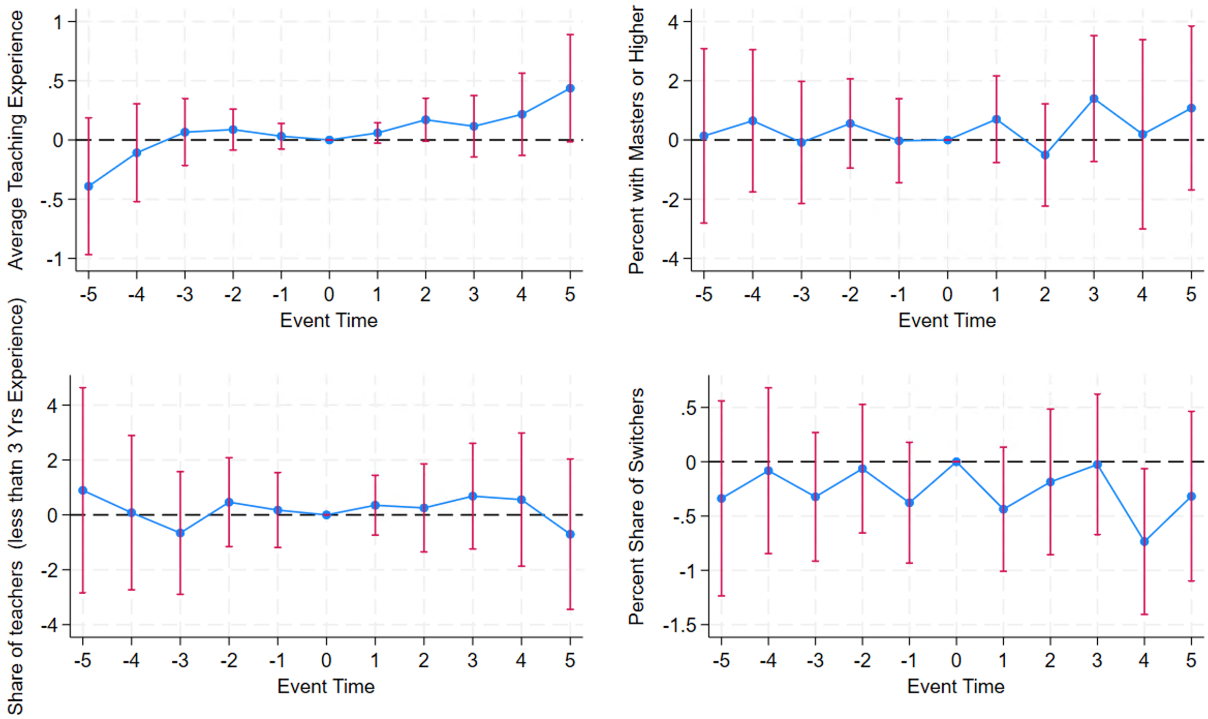

Note: Point estimates and 95% confidence intervals of parameters $\beta_t D_{it}$ in Equation (2). Average experience, total teachers with a Master's Degree or higher, and total new teachers (< 3 years of experience) calculated from aggregating teacher-level data up to the school level. Treatment is continuous. The teacher dataset used incorporates NLP and Levenshtein distance.

From Figure 11, we rule out significant compositional changes in the ELA and Math teachers at high schools. The average level of teaching experience remains largely unchanged, although there is a slight, statistically insignificant upward trend. The total number of teachers holding a Master's degree or higher remains stable, as does the proportion of newly hired teachers in a school. Furthermore, we find no evidence of an increased share of teachers that "switched" in treated schools. Given that the share of switchers does not increase but the average teaching experience shows an upward, albeit statistically insignificant, trend, it could be that teachers in treated schools are staying longer, resulting in lower teacher turnover rates.

10.1 Taking Stock

From our results in Figures 10 and 11, we rule out the plausible mechanism that improvements in student outcomes were driven by "higher quality" teachers relocating to better-paying districts. Consequently, we conclude that the dynamics of the teacher labor market in New Jersey differed from those observed in Wisconsin. A possible reason why we do not observe this effect could be the nature of the salary increase. Since teachers are not paid based on performance and many school districts in NJ eventually adopted these salary schedules, there may be fewer incentives for teachers to switch school districts.

On that note, although we rule out switching effects, we cannot causally argue for or against the motivation mechanism. While it appears that improvements in student outcomes could have been driven by improved teacher morale, we do not have a direct way to measure this, given we lack teacher-linked student outcome data.[17] Another plausible mechanism is that while both treated and not-yet-treated schools are hiring new teachers at the same rate before and after the act, "treated" school districts are hiring higher-quality individuals into the teaching profession—i.e., new, higher-quality teachers are deciding to work in these "treated" schools.

11 Conclusion

This paper examines the impact of raising teacher salaries on student outcomes by exploiting the variation in the years districts approved collective bargaining agreements that set minimum teacher salaries between \$50,000 and \$54,000. Using a staggered Difference-in-Differences (DiD) approach with a continuous treatment framework, we show that salary increases led to significant improvements in 4th-grade and high school Math and ELA scores. We also observe modest gains in graduation rates and college enrollment. These findings underscore the crucial role teacher compensation plays in shaping educational success.

Importantly, we show that the financial resources required to fund the salary increases did not come from reallocating funds away from more experienced teachers. Instead, we show that districts achieved this by hiring fewer staff members and shifting the amount saved towards boosting compensation. Our analysis reveals that the differences in per-pupil spending remained constant, and thus we rule out any confounding effects due to changes in non-instructional spending. We also show that salary increases for non-teaching staff remained unaffected. This lends greater credibility to the conclusion that the observed improvements in student performance are attributable to improved staff compensation rather than broader budgetary shifts.

The contributions of this study are primarily twofold. First, we advance the relatively sparse literature on the causal impact of teacher salary increases by demonstrating a clear and robust link between higher pay and improved student outcomes. The improvements in outcomes we observe—an increase in proficiency scores in the range of 0.05 to 0.07 standard deviations—are comparable to, and in some cases exceed, previous estimates (Section 9). Second, we show that the observed gains are not due to districts hiring "better quality" teachers from other schools. Instead, the results are likely driven by a combination of other factors such as schools hiring more qualified new teachers and improved teacher morale among existing staff. Given our data, we are unable to precisely determine which of these mechanisms are in play. Future research should thus focus on exploring these channels as potential mechanisms.

While our findings indicate that increasing teacher salaries can be a highly effective strategy for improving student outcomes, it is important to acknowledge certain limitations

[17] We have also submitted a data request for teacher-linked student performance data, which is currently under review.

of our study. First, aggregate student scores, particularly those used in this analysis, are an imperfect measure of student success. Ideally, a value-added model that utilizes classroom-level data linked directly to individual teachers would provide a more precise estimate of the impact of salary increases. We hope to acquire such data in future studies for a more accurate assessment. Furthermore, even with teacher-linked student data, the most robust analysis would involve evaluating long-term student outcomes—such as college completion and income trajectories. Regarding the generalizability of the study, it is worth noting that the average salaries of New Jersey public school teachers are significantly higher than the national average. Therefore, we posit that increasing teacher salaries in other states would yield similar, if not greater, improvements in student performance.

In conclusion, our findings provide compelling evidence that raising teacher salaries has a substantial and positive impact on student outcomes. Policymakers should seriously consider the broader implications of these findings when designing educational policies—particularly in regions where teacher compensation has historically been insufficient. By demonstrating that even modest salary increases can lead to significant improvements in student performance, this study shows that prioritizing teacher pay can enhance educational quality and, ultimately, improve the life prospects of students. All in all, lifting teacher salaries for all teachers—regardless of whether they are high performing—appears to improve student outcomes.

References

Angrist, J. D., & Lavy, V. (1999). Using maimonides' rule to estimate the effect of class size on scholastic achievement. *The Quarterly Journal of Economics*, *114*(2), 533–575.

Baron, E. J. (2018). The effect of teachers' unions on student achievement in the short run: Evidence from wisconsin's act 10. *Economics of Education Review*, *67*, 40–57.

Biasi, B. (2021). The labor market for teachers under different pay schemes. *American Economic Journal: Economic Policy*, *13*(3), 63–102. https://doi.org/10.1257/pol.20200295

Biasi, B., Fu, C., & Stromme, J. (2021). *Equilibrium in the market for public school teachers: District wage strategies and teacher comparative advantage* (tech. rep. No. w28530). National Bureau of Economic Research.

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200–230. https://doi.org/10.1016/j.jeconom.2020.12.001

De Chaisemartin, C., & d'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American economic review*, *110*(9), 2964–2996.

De Chaisemartin, C., & d'Haultfoeuille, X. (2024). Difference-in-differences estimators of intertemporal treatment effects. *Review of Economics and Statistics*, 1–45.

Dube, A., Girardi, D., Jordà, Ò., & Taylor, A. M. (2023). *A local projections approach to difference-in-differences* (NBER Working Paper No. 31184). National Bureau of Economic Research. https://doi.org/10.3386/w31184

Eberts, R. W., & Stone, J. A. (1988). Student achievement in public schools: Do principals make a difference? *Economics of Education Review*, *7*(3), 291–299. https://doi.org/10.1016/0272-7757(88)90019-3

Fryer, R. G. (2013). Teacher incentives and student achievement: Evidence from new york city public schools. *Journal of Labor Economics*, *31*(2), 373–407. https://doi.org/10.1086/667757

Goodman, S. F., & Turner, L. J. (2013). The design of teacher incentive pay and educational outcomes: Evidence from the new york city bonus program. *Journal of Labor Economics*, *31*(2), 409–420. https://doi.org/10.1086/667846

Han, E., & Garcia, E. (2022). Teachers' base salary and districts' academic performance: Evidence from national data. *Sage Open*. https://doi.org/10.1177/21582440221082138

Hanushek, E. A. (1997). Assessing the effects of school resources on student performance: An update. *Educational Evaluation and Policy Analysis*, *19*(2), 141–164. https://doi.org/10.3102/01623737019002141

Hanushek, E. A. (2003). The failure of input-based schooling policies. *The Economic Journal*, *113*(485), F64–F98. https://doi.org/10.1111/1468-0297.00099

Hanushek, E. A. (2015). *Education production functions*. Springer. https://doi.org/10.1007/978-1-4419-1428-6_423

Hanushek, E. A., Piopiunik, M., & Wiederhold, S. (2019). The value of smarter teachers: International evidence on teacher cognitive skills and student performance. *Journal of Human Resources*, *54*(4), 857–899. https://doi.org/10.3368/jhr.54.4.0617.8961R1

Henry, G. T., & Redding, C. (2020). The consequences of leaving school early: The effects of within-year and end-of-year teacher turnover. *Education Finance and Policy*, *15*(2), 332–356. https://doi.org/10.1162/edfp_a_00274

Jackson, C. K., Johnson, R. C., & Persico, C. (2016). The effects of school spending on educational and economic outcomes: Evidence from school finance reforms. *The Quarterly Journal of Economics*, *131*(1), 157–218. https://doi.org/10.1093/qje/qjv036

Jackson, C. K., & Mackevicius, C. L. (2024). What impacts can we expect from school spending policy? evidence from evaluations in the united states. *American Economic Journal: Applied Economics*, *16*(1), 412–446. https://doi.org/10.1257/app.20220279

Ladd, H. F., & Sorensen, L. C. (2017). Returns to teacher experience: Student achievement and motivation in middle school. *Education Finance and Policy*, *12*(2), 241–279. https://doi.org/10.1162/EDFP_a_00194

Miller, A. (2013). Principal turnover and student achievement. *Economics of Education Review*, *36*, 60–72. https://doi.org/10.1016/j.econedurev.2013.05.004

Muralidharan, K., Pradhan, M., & Rogers, H. (2016). Double for nothing? https://www.nber.org/papers/w21806

Muralidharan, K., & Sundararaman, V. (2011). Teacher performance pay: Experimental evidence from india. *Journal of Political Economy*, *119*(1), 39–77.

New Jersey Department of Education. (2024). School performance reports [Accessed: 2024-08-23]. https://www.nj.gov/education/doedata/fact.shtml

New Jersey Department of the Treasury. (2011). Pension reform 2011 [[Accessed: 2024-08-05]].

Ng, K. (2024). The effects of teacher tenure on productivity and selection. *Economics of Education Review*, *101*, 102558.

Rambachan, A., & Roth, J. (2020). An honest approach to parallel trends. *Working Paper*. https://github.com/asheshrambachan/HonestDiD

Springer, M. G., Ballou, D., Hamilton, L., Le, V.-N., Lockwood, J. R., McCaffrey, D. F., Pepper, M., & Stecher, B. M. (2012). Teacher pay for performance: Experimental evidence from the project on incentives in teaching (point). *Economics of Education Review*, *31*(1), 1–14. https://doi.org/10.1016/j.econedurev.2011.10.003

A Appendix

Table 8: Year Negotiation was Approved

Year	**Frequency**	**Percent**	**Cumulative Percent**
Never Treated	39	2.31%	2.31%
2006	9	0.53%	2.85%
2007	14	0.83%	3.68%
2008	64	3.80%	7.47%
2009	144	8.54%	16.01%
2010	204	12.10%	28.11%
2011	248	14.71%	42.82%
2012	157	9.31%	52.14%
2013	149	8.84%	60.97%
2014	92	5.46%	66.43%
2015	101	5.99%	72.42%
2016	86	5.10%	77.52%
2017	105	6.23%	83.75%
2018	80	4.74%	88.49%
2019	85	5.04%	93.53%
2020	44	2.61%	96.14%
2021	39	2.31%	98.46%
2022	4	0.24%	98.70%
2023	21	1.25%	99.94%
2024	1	0.06%	100.00%

Note: Since our school-level outcomes are from 2003-2019, we have a decent number of schools acting as controls. Never treated indicates that 39 schools have yet to adopt this policy. Year 2006 denotes that the negotiation was approved to go into effect in the 2006-2007 school year.

Table 9: Difference-in-Difference Estimates (Event Study)

Variable	Coefficient	Std. Err.	Lower Bound	Upper Bound
Pre Policy	-322.85	173.92	-663.74	18.03
Post Policy	1234.18	470.56	311.91	2156.46
tm4	-623.83	210.80	-1036.99	-210.68
tm3	-282.90	210.23	-694.95	129.15
tm2	-38.20	218.41	-466.27	389.87
tm1	-346.47	188.86	-716.63	23.70
tp1	474.83	150.16	180.52	769.14
tp2	672.47	268.66	145.91	1199.03
tp3	909.39	394.03	137.10	1681.68
tp4	998.36	530.15	-40.71	2037.44
tp5	1622.20	546.57	550.95	2693.45
tp6	1564.13	586.87	413.89	2714.37
tp7	1777.37	695.86	413.51	3141.23
tp8	1854.71	934.52	23.08	3686.34

Note: Table reports point estimates of Salary increases with standard errors clustered at the district level. Lower bound and upper bound represent the 95% confidence interval. tm denotes pre-policy, and tp denotes post-policy. The sample includes all full-time staff (both teaching and non-teaching).

Compensation Glossary

- **Average Salary** – the base salary cost divided by the total number of full-time employees (FTE) on the scattergram.

- **Base Salary Cost** – The total of each step on the guide multiplied by each corresponding step on the scattergram. Other amounts may or may not be included, such as longevity, ratio differentials, extra-curricular activities, stipends, black seal amounts, building stipends, etc.

- **Breakage** – The amount of dollars saved between the salary of a departing employee (retirement, resignation, and leave of absence) and the new employee who is replacing the departed employee.

  Example:

  - $50,000 salary of retiring employee
  - $30,000 replacement employee
  - $20,000 breakage

- **Bubble/Balloon** – An abnormal separation between two steps on a salary guide.

  Example:

  - Step 13 - $39,000
  - Step 14 - $40,000
  - Maximum - $50,000
  - Increment - $10,000 or 25%

- **Cumulative Earnings** – The total sum of all salaries in a specified time period or career. NJEA Research calculates the 10-, 20-, and 30-year earnings based on a long-standing formula of 5 years on the BA column and the remaining years on the MA column. Longevity is added, as are any other negotiated amounts at the appropriate time.

- **Guide Movement** – A movement from one step on a guide to the next higher step on that guide. Horizontal movement would be movement to a higher credit/degree/level guide based on a specified criteria.

- **Horizontal/Lane/Column** – A specific list of salaries with a minimum, maximum, and number of steps.

- **Increment** – The dollar amount of the salary increase an individual receives when they advance a step on the guide.

  Example:

  - Step 1 - $50,000
  - Step 2 - $51,000
  - Increment - $1,000

- **Increment Cost** – The dollar amount of the increment multiplied by the number of individuals that will receive that increment for a contract year.

- **Longevity** – Additional money paid to an employee above the salary guide. It is usually based on years of service to either the school district or the profession in general. It is usually a specified dollar amount, but can also be a percent of salary.

  Example:

  - $1000 additional for 15 years of service to the district
  - or 3% of individual salary for 15 years of service to the district.

- **Maximum** – The highest step on the published salary guide. It may also be called the career rate.

- **Minimum** – The beginning step of a guide that is considered to be the hiring step with no experience.

- **Off Guide Salaries** – Additional salaries that are paid above the printed salary guides. They are actually additional steps on a guide.

- **Salary Guide** – A chart that shows the dollar value of each step and level/category.
- **Scattergram** – A chart showing the number of employees on each step and level/category of a salary guide. These employees will generally be in the full-time equivalency (FTE) category of employment.

A.1 Teacher Mobility and Inequality: Theoretical Model

This model examines how wage differentials between school districts lead to sorting within a fixed pool of teachers, exacerbating inequalities in teacher quality and student outcomes.

Assumptions

- **No significant increase in Teacher Pool:** The total number of teachers N is relatively constant i.e. individuals who did not want to be teachers do not get into teaching due to the increase in pay.
- **Teachers:** Each teacher i has a quality level $q_i \sim F(q)$, where $F(q)$ is the distribution of teacher quality with mean μ_q and variance σ_q^2.
- **Districts:** Each district j offers a wage w_j, and districts differ in student composition, represented by λ_j, the fraction of low-achieving students.
- **Utility Function:** The utility of teacher i in district j is:

  $$U_i^j = w_j + \phi(\lambda_j) + \epsilon_i^j$$

  where $\phi(\lambda_j)$ represents non-pecuniary disutility from teaching in low-performing districts, and ϵ_i^j is a teacher-specific preference shock.

Districts' Problem: Setting Wages

Each district j maximizes the average quality of its teachers q_j by offering competitive wages, subject to its budget constraint:

$$\max_{w_j} q_j = E[q_i \mid U_i^j \geq U_i^k \quad \forall k]$$

subject to:

$$w_j N_j \leq B_j$$

where N_j is the number of teachers and B_j is the district's budget.

Teachers' Decision: Sorting Based on Utility

Teachers sort themselves across districts to maximize their utility. Teacher i will choose district j over district k if:

$$U_i^j = w_j + \phi(\lambda_j) \geq w_k + \phi(\lambda_k)$$

Rearranging:

$$w_j - w_k \geq \phi(\lambda_k) - \phi(\lambda_j)$$

This shows that a wage differential $w_j - w_k$ must be large enough to compensate for differences in student composition.

Wage Differentials and Teacher Quality

Given the fixed teacher pool, the average quality of teachers in district j is:

$$q_j = E[q_i \mid U_i^j \geq U_i^k \quad \forall k]$$

High-wage districts attract better teachers, so $q_H > q_L$, where q_H and q_L represent the average quality of teachers in high- and low-wage districts, respectively.

Expanding Inequality (with Fixed Teacher Pool)

Wage differentials lead to growing inequality in teacher quality:

$$\Delta q = q_H - q_L$$

As wage differentials $w_H - w_L$ increase, this gap expands over time:

$$\Delta q_t = \Delta q_{t-1} + f(w_H - w_L)$$

Even though N is fixed, wage competition reallocates teachers, causing the inequality in teacher quality and student outcomes to persist:

$$\lim_{t \to \infty} \Delta q_t > 0$$

B Event Study: Alternate Specification

Figure 12: Same Switchers (Y-axis is Percentage of students meeting Proficiency)

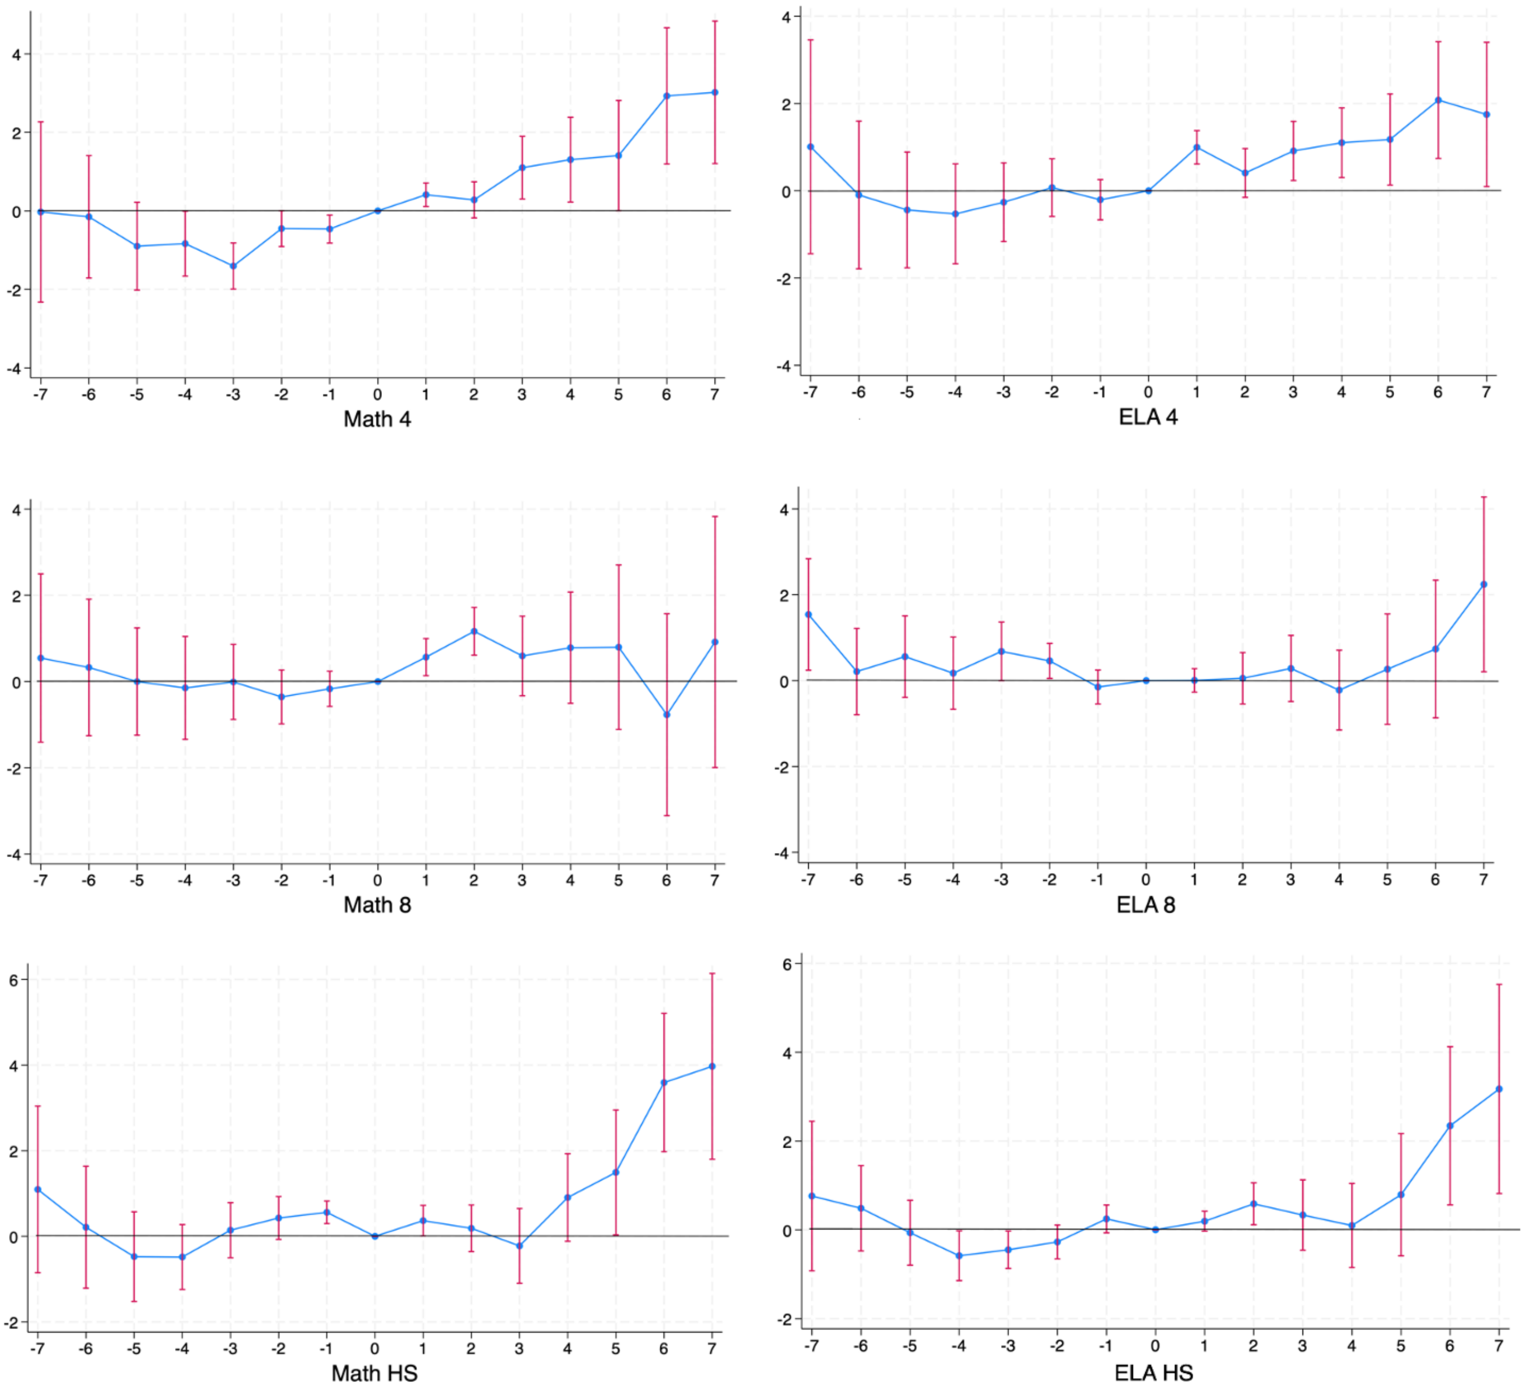

Note: Point estimates and 95% confidence interval of parameters $\beta_t D_{it}$ in Equation 2 at the school level. Y-axis shows changes in percentage of students meeting proficiency. X-axis refers to time period before and after the treatment. Reference year is time period 0, which indicates the year treatment took place. Treatment is non-binary and run on a strongly balanced panel(same-switchers). Although not shown here, effects hold for graduation rates as well. We cannot run a same switcher model estimate on percentage of students enrolled in college due to missing data for years 2011 and 2012.

Figure 13: Standardized Scores

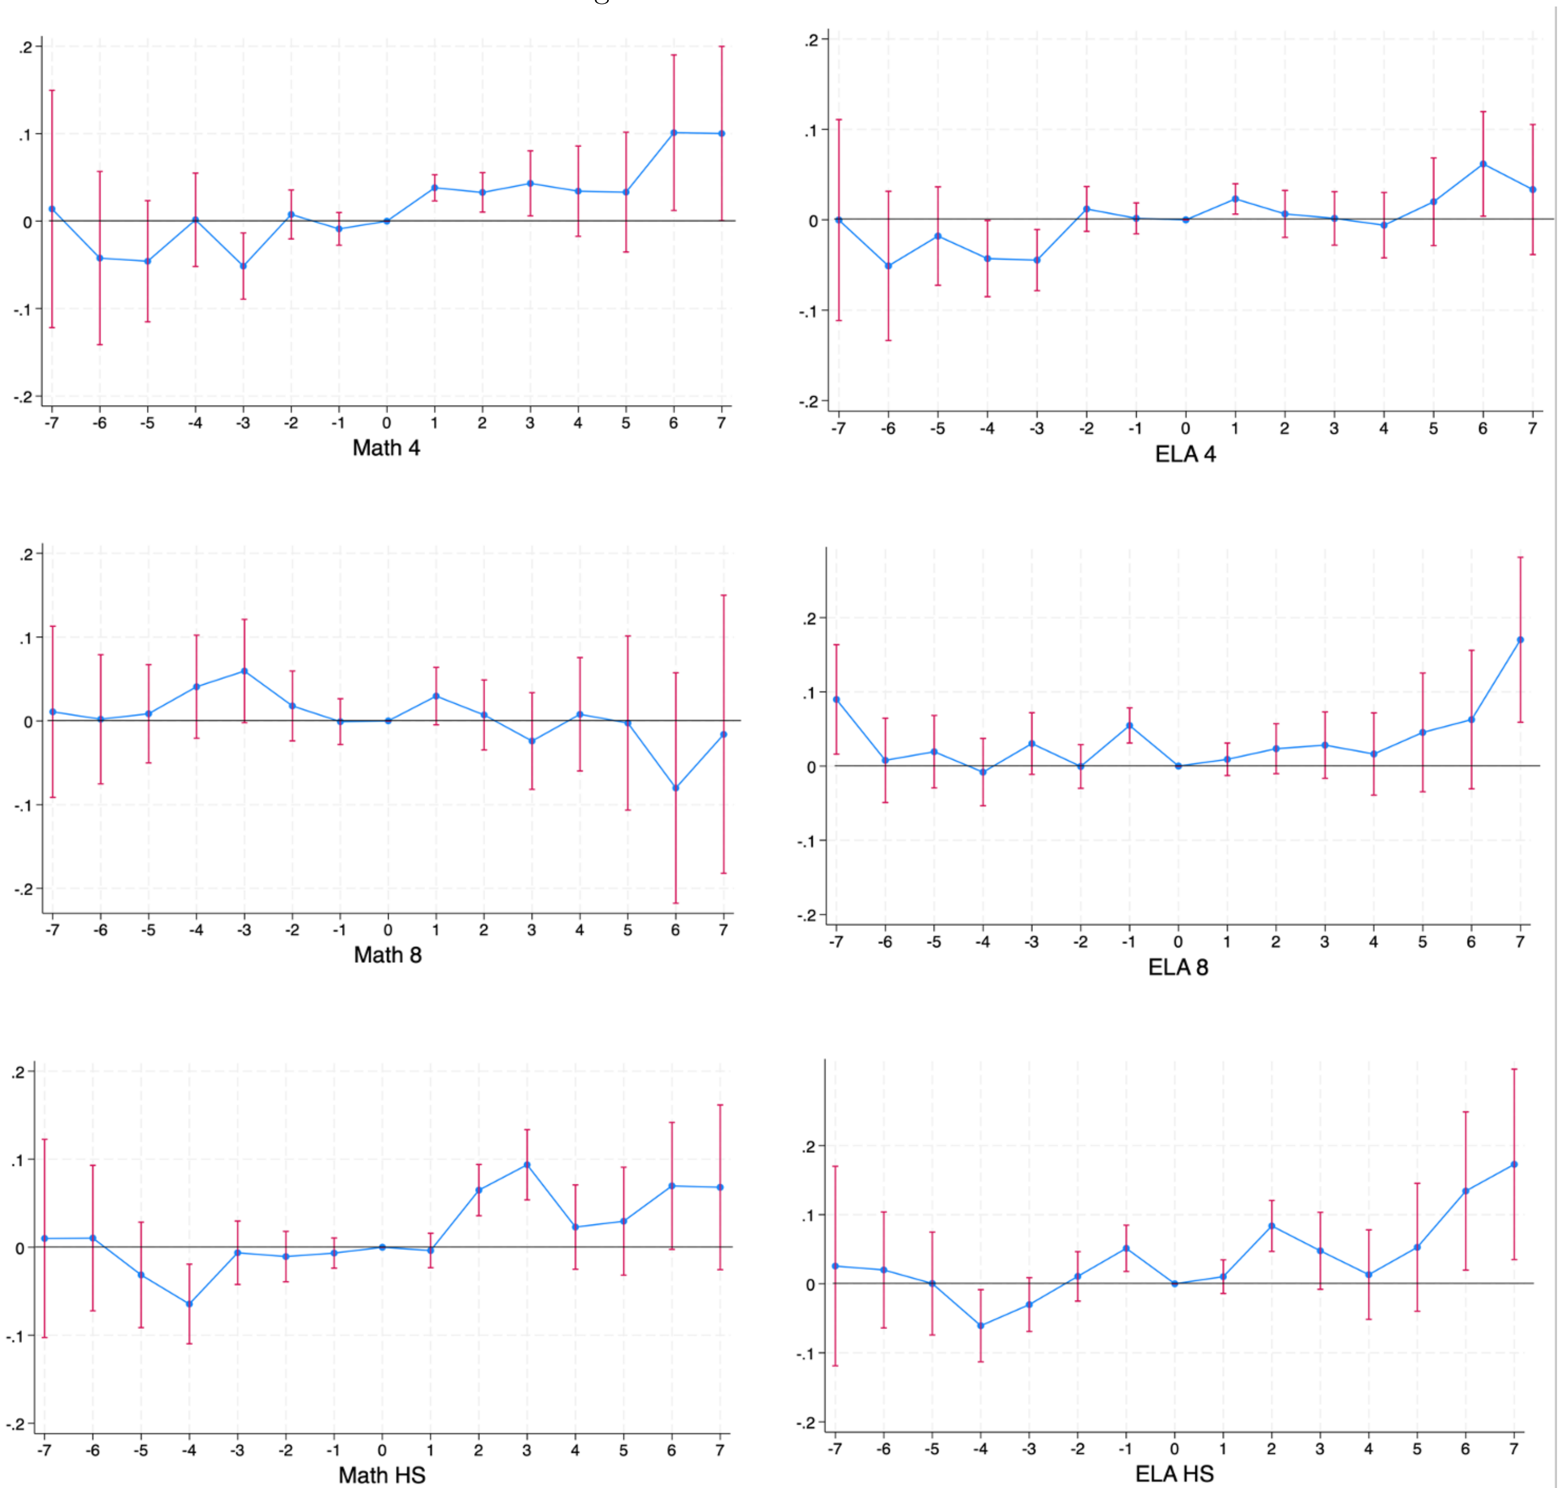

Note: Point estimates and 95% confidence interval of parameters $\beta_t D_{it}$ in Equation 2 at the school level. Y-axis shows changes in standardized values of percentage of students meeting proficiency. Each outcome is standardized by year. X-axis refers to time period before and after the treatment.

C Event Study using Alternate Estimators

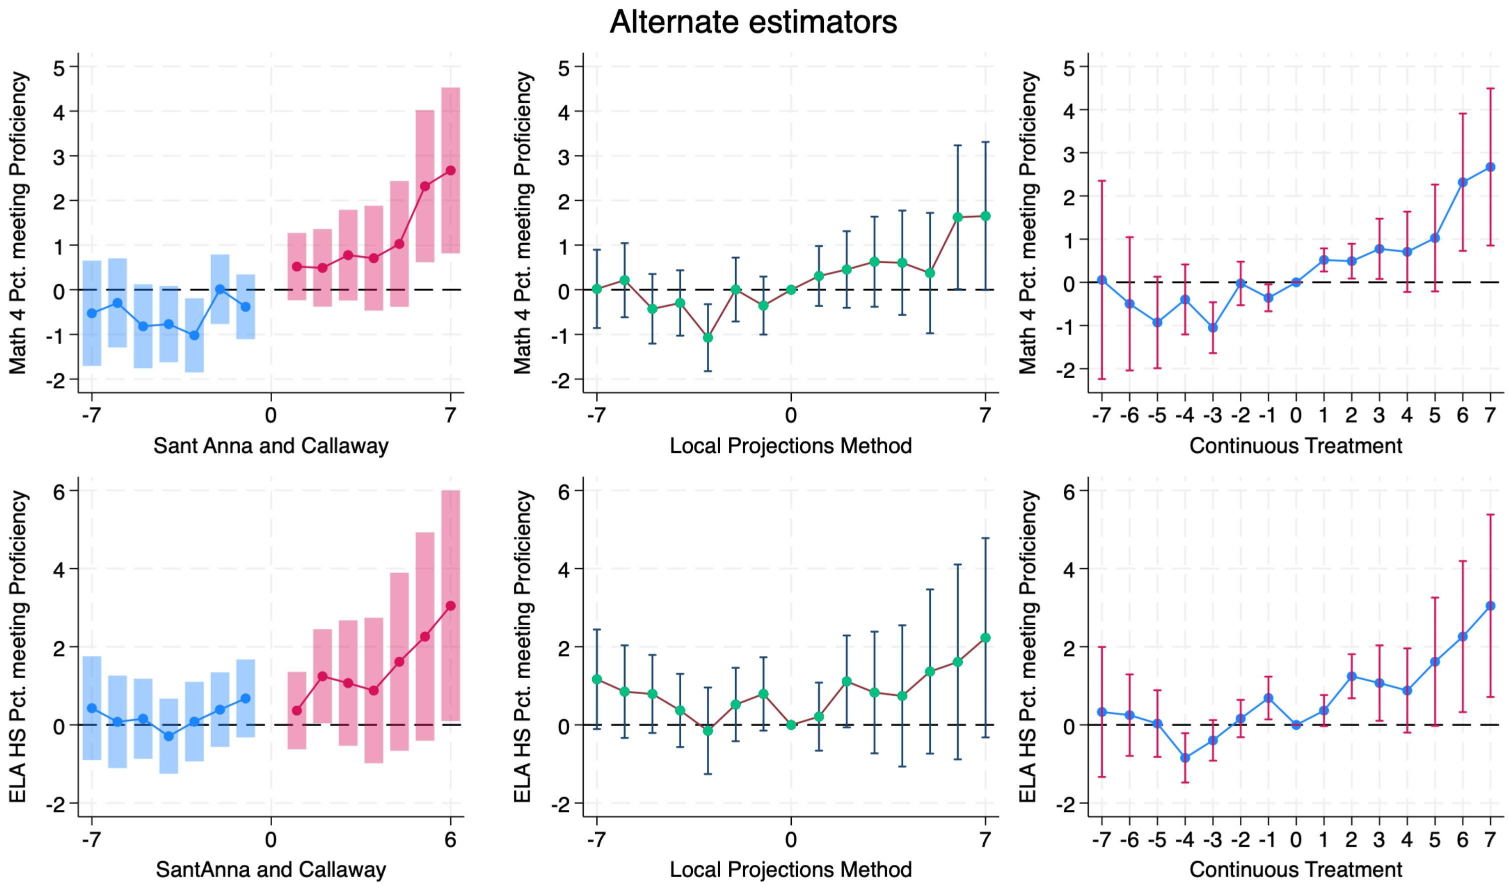

Note: Point estimates and 95% confidence interval of parameters $\beta_t D_{it}$ in Equation (2). The first columns shows event study using Callaway and SantAnna DiD estimators. Pink shaded areas denote post treatment periods and blue shaded areas denote pre treatment periods. The white gap is the reference period. The second column shows event study using the Local projections DiD. The third column showns event study using continuous treatment. This is the same event study we show in our results of the paper. Included here for reference. Although we only show results for Math 4 and high school ELA, results hold for all other grades suggesting that our results are robust to using alternate estimators. Standard errors are the largest when using SantAnna and Callaway.

D Event Study controlling for socio-economic status

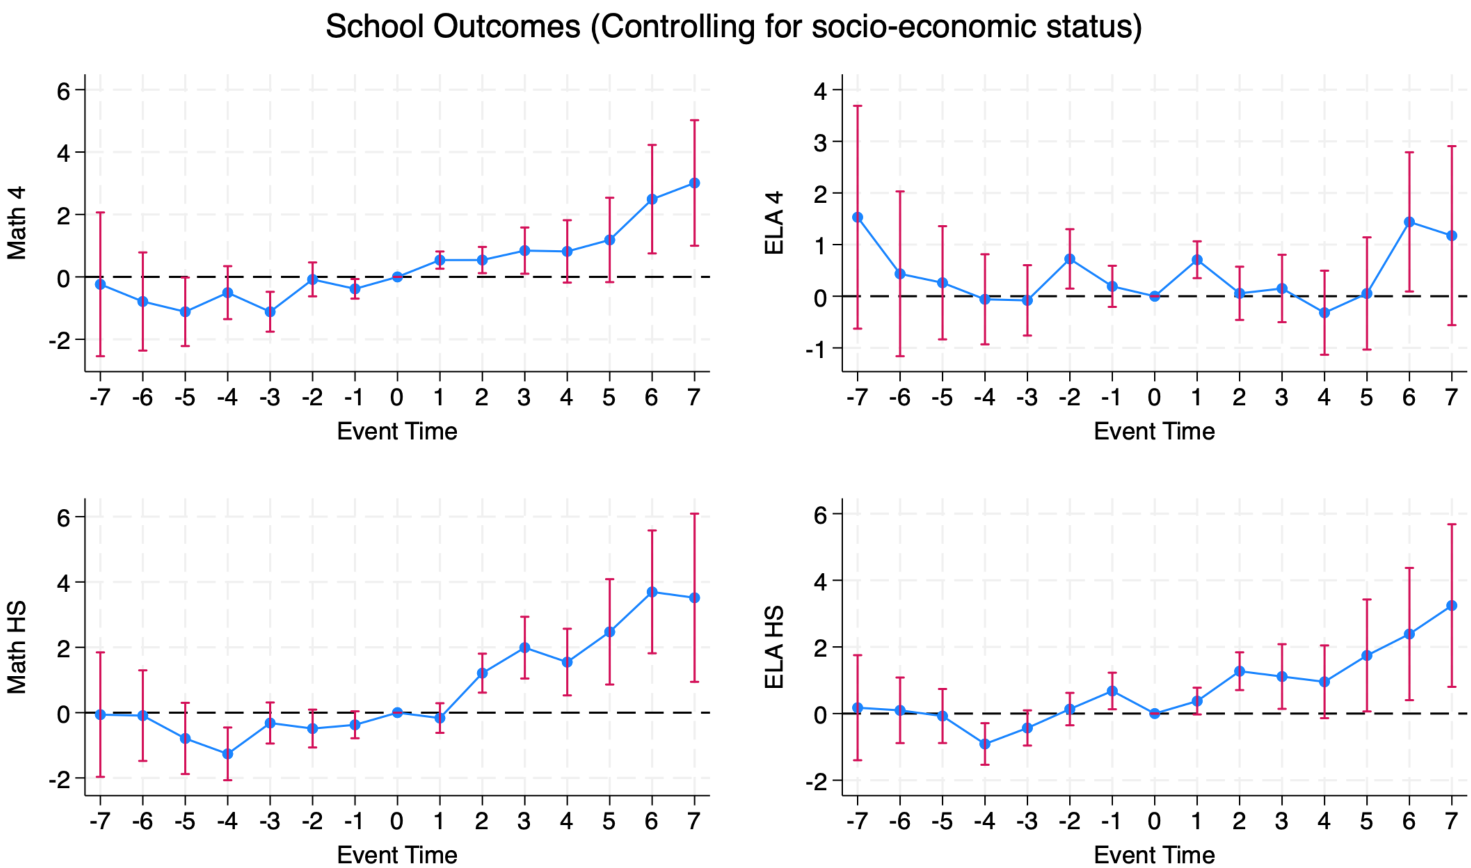

Note: Point estimates and 95% confidence interval of parameters $\beta_t D_{it}$ in Equation (2) controlling for socio-economic status based on district factor group. The classification is done based on Table 10.

Table 10: Distribution of classifications across categories

Socio-economic Classification	**A**	**B**	**CD**	**DE**	**FG**	**GH**	**I**	**J**	**Total**
1	3,927	3,230	2,669	0	0	0	0	0	9,826
2	0	0	0	3,927	3,995	4,284	0	0	12,206
3	0	0	0	0	0	0	5,457	1,156	6,613
Total	3,927	3,230	2,669	3,927	3,995	4,284	5,457	1,156	28,645

From lowest socioeconomic status to highest, the categories are A, B, CD, DE, FG, GH, I, and J. Number here denotes total observations across all years for all schools.

E Event Study: Random Treatment Year

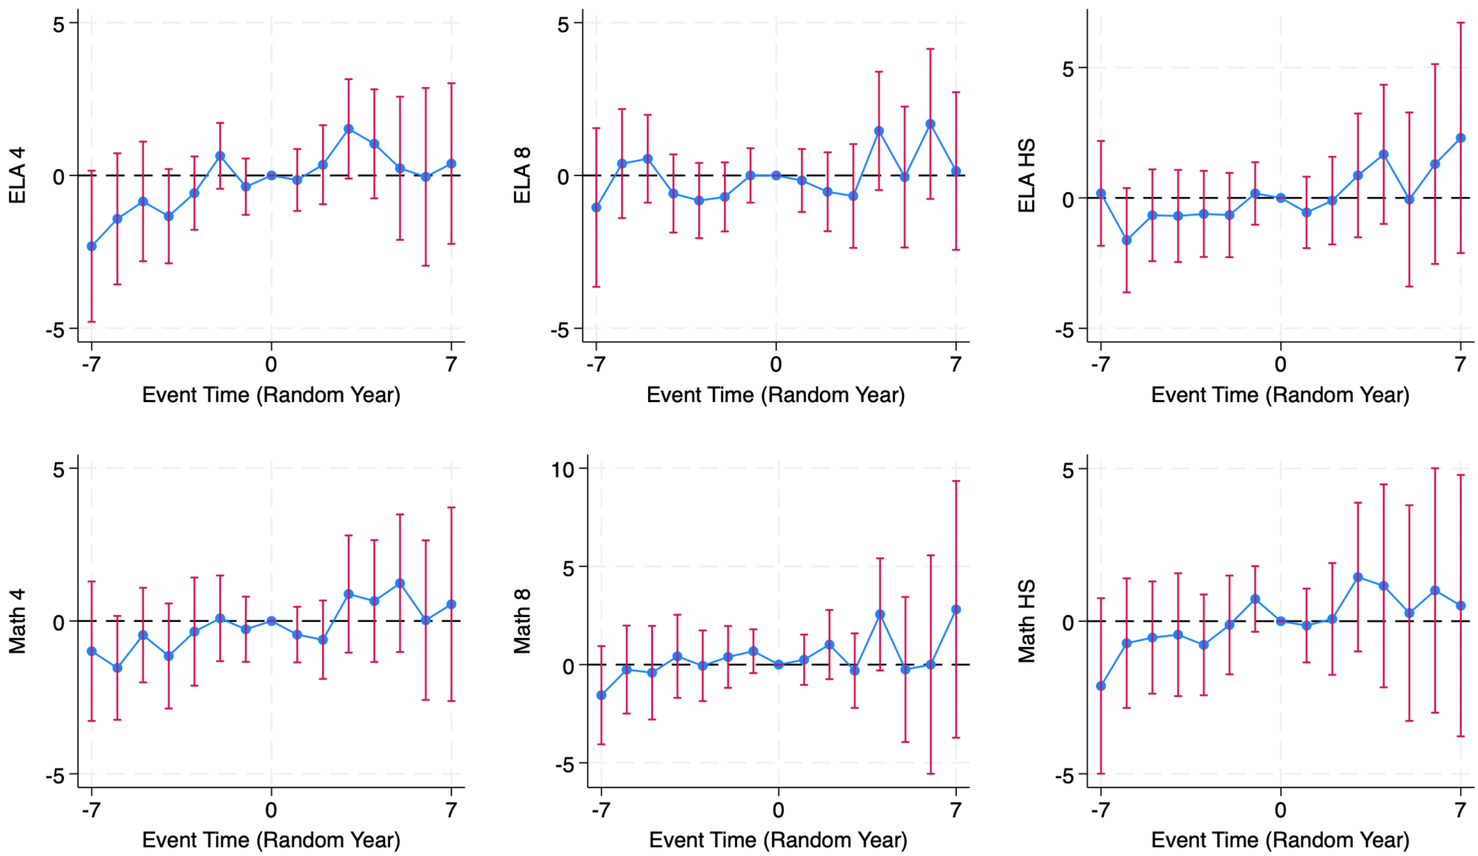

Note: Point estimates and 95% confidence interval of parameters $\beta_t D_{it}$ in Equation (2). Treatment year three years from the actual year the negotiation was passed. Event study shows no positive effects when the year of treatment is randomized.