ANNENBERG
BROWN UNIVERSITY

EdWorkingPaper No. 22-512

Pension Reform and Labor Supply

Andrew C. Johnston
University of California at Merced

Jonah Rockoff
Columbia University

As unfunded pension liabilities grow, governments experiment with ways to curb costs. We examine the effect of a representative cost-cutting reform on the retention and productivity of workers. The reform reduced pension annuities and increased penalties for early retirement, projected to save 8 percent of revenues. We leverage administrative records and a discontinuity in the reform to estimate its effect on labor supply. The reform slightly increased worker retention, and we can rule out small attrition effects. The reform had no effect on worker output. The extensive and intensive margins of labor supply appear to be maintained under the reform.

VERSION: January 2022

Suggested citation: Johnston, Andrew C., and Jonah Rockoff. (2022). Pension Reform and Labor Supply. (EdWorkingPaper: 22-512). Retrieved from Annenberg Institute at Brown University: https://doi.org/10.26300/5pat-r352

Pension Reform and Labor Supply

Andrew C. Johnston[*] Jonah Rockoff[†]

January 4, 2022

Abstract

As unfunded pension liabilities grow, governments experiment with ways to curb costs. We examine the effect of a representative cost-cutting reform on the retention and productivity of workers. The reform reduced pension annuities and increased penalties for early retirement, projected to save 8 percent of revenues. We leverage administrative records and a discontinuity in the reform to estimate its effect on labor supply. The reform slightly increased worker retention, and we can rule out small attrition effects. The reform had no effect on worker output. The extensive and intensive margins of labor supply appear to be maintained under the reform.

JEL: H75, J26, M52

[*]University of California at Merced, 5200 Lake Road, Merced CA, 95340. Email: acjohnston@ucmerced.edu. The authors are grateful to the Texas Education Agency for providing the data, and we thank Chad Aldeman, Adam Bloomfield, Mark Duggan, Maria Fitzpatrick, Dillon Fuchsman, Todd Gormley, Robert Jensen, Cory Koedel, Alex Mas, Olivia S. Mitchell, Josh McGee, Joshua Rauh, Laura Quinby, and Gema Zamarro for constructive comments. Ryan Santos provided outstanding research assistance

[†]Columbia Business School, 3022 Broadway, New York NY, 10027. Email: jonah.rockoff@columbia.edu

1 Introduction

The cost of pensions for public servants looms large across the modern world. In some countries, unfunded pension liabilities are larger than official government debt (Hanif et al., 2016; Mitchell, 2020). The World Economic Forum (2017) projects that unfunded pension liabilities "will grow by 5 percent each year to $400 trillion by 2050...an additional $28 billion of deficit each day." These costs have pressed state and local governments to cut education and infrastructure investments, possibly hampering economic growth and mobility (Barro, 1991; Duflo, 2001; Czernich et al., 2011; Chetty et al., 2014; Anzia, 2017; Koedel, 2019). Alarmed, policymakers have attempted to curb pension costs with tailored reforms (e.g., Fitzpatrick, 2017).

In this paper, we use administrative staffing data and a discontinuity to evaluate one such reform that affected the sector containing the majority of public-sector workers—those employed by public schools. The reform cut costs on two fronts. First, it reduced the value of annuities workers would receive in retirement by 3–11 percent, depending on a worker's earnings history. Second, the reform multiplied the penalties for retiring early by a factor of three. Before the reform, a worker could retire at age 55 with 90 percent of the benefits she would have received under the full-retirement formula. After the reform, retiring at 55 meant the annuity was cut in half. The state's congressional budget office estimated the reform would save the state $250 million each year, approximately 8 percent of yearly pension contributions (Legislative Budget Board, 2005).

Opponents argued that such cuts undermine public education by reducing the appeal of working for public schools, decreasing retention, and eroding the quality of public education (Sabatosa, 2021). Under economic theory, the net effect of the policy on retention is somewhat more ambiguous. On one side of the ledger, the continuation value of retention is diminished by decreasing annuities in retirement, reducing retention. On the other side, early retirement is made less attractive so workers may be more likely to persist to full retirement.

To evaluate the effect of the reform, we leverage the fact that workers beyond a threshold were grandfathered into the pre-existing pension regime while those who didn't meet the threshold were placed in the new regime. The cutoff conduces a discontinuity-style quasi-experimental design. A worker was grandfathered if, in the fall of 2005, she was at least 50 years of age, had at least 25 years of experience, or the sum of her age and experience exceeded 70. By comparing similar workers on either side of the threshold, we find that the reform increased retention rates by 1–2 percentage points over 15 years.

We examine whether the reform affected productivity in public schools, as proxied by gains in student achievement. In some models, the reform reduces output by depressing worker morale (Weakliem and Frenkel, 2006; Kube et al., 2013). In others, the reform elicits greater effort by raising the cost of premature dismissal—one of the theoretical rationales for deferred compensation vehicles like pensions (Lazear, 1979; Gustman et al., 1994). We estimate the effect of the reform using a difference-in-differences approach. The identification strategy leverages the timing of the reform (the first difference) and fact that some school-grade cells are more exposed to the reform than others, based on the age and experience composition of teachers in each cell (the second difference). In our base specification the reform is associated with achievement gains in math and reading, but a specification that uses across grade variation *within* a school finds the reform has no significant effect on productivity.

The 2005 Texas reform was a harbinger of a wave of similar reforms across the country. Between 2009 and 2014, three quarters of states implemented cost-cutting reforms, most of which shared the elements of the reform we study in Texas. And since 2009, every state in the union has passed cost-saving measures for their state pension programs (Aubry and Crawford, 2017; Quinby et al., 2018). Like Texas, nearly two-thirds of states reduced pension annuities by calculating final average salary over a longer period, almost always expanding the base to five years from three;[1] and approximately 40 percent of states cut early-retirement generosity, similar to the Texas reform (Brainard and Brown, 2018).[2]

[1]These include Alabama, Arizona, Arkansas, California, Colorado, Connecticut, Florida, Hawaii, Illinois, Iowa, Kansas, Louisiana, Maryland, Massachusetts, Montana, Nebraska, Nevada, New Hampshire, New Jersey, New Mexico, New York, Ohio, Oklahoma, Pennsylvania, Rhode Island, South Carolina, another reform in Texas for public workers, Utah, Virginia, West Virginia, and Wyoming.

[2]This includes California, Connecticut, Delaware, Illinois, Iowa, Kentucky, Maine, Maryland, Massachusetts, Minnesota, Missouri, Montana, Nevada, New York, North Dakota, South Carolina, Ohio, Wash-

The reform we examine has a few features useful for analysis. First, the cutoffs allow us to compare similar workers operating in the same work setting while under different pension regimes. This provides for a careful comparison of like employees treated differently. Second, because the Texas reform was relatively early compared to the reforms of other states, it allows us to study the effect over a long window to measure its consequence. Third, and especially important, the reform affected *current* workers. The vast majority of pension reforms only affect new workers, which complicates analysis.[3] If a reform grandfathers all incumbents (as most do), the treatment is confounded with selection on entry, making it difficult to isolate the impact of the new policy on either retention or output. Since the reform we examine affected existing teachers, we can isolate the effect of the reform on behavior from the influence of selection.

This paper contributes to a literature examining the influence of pensions on worker turnover and performance (Friedberg and Webb, 2005; Costrell and McGee, 2010; Goldhaber et al., 2017; Fitzpatrick, 2019; Ni et al., 2020).[4] Brown (2013) uses a differencing-and-bunching approach to evaluate the effect of pensions on labor supply. Manoli and Weber (2016) use retirement-eligibility notches in Austria for the same purpose. Koedel and Xiang (2017) and Ni et al. (2021) examine the effect of a local pension enhancement on retention. Whereas past work examined the retention effects of the status quo or enhancements, what is unique about our setting is our ability to examine the effect of a cost-cutting reform on both the retention and output of workers. Cutoffs avail us of simple treatment-control comparison where like workers in the same employment setting make choices under different pension regimes. The environment also provides a rare opportunity to examine the consequences of pensions on worker output which isn't normally observable in other settings. We can thus cast light on the reform's influence on both the extensive and intensive margins of labor supply.

ington, Wyoming.

[3] For instance, almost every state since 2007 has altered the pension parameters for new hires while preserving benefits for existing workers (including California, Massachusetts, Florida, Colorado, Alabama, Alaska, Arizona, Arkansas, and Connecticut) (Brainard and Brown, 2018). Over half of all pension programs in the country raised retirement ages or vesting requirements. The vast majority only affected new hires (Quinby et al., 2018).

[4] A related literature carefully examines how workers value their retirement benefits (Fitzpatrick, 2015; Biasi, 2019; Fuchsman et al., 2020; Ni and Podgursky, 2016). See also Johnston (2021)

2 The Pension Reform

The Texas reform bill was introduced in March of 2005 by bi-partisan coauthors, Craig Eiland (a House Democrat from Galveston) and Robert Duncan (a Senate Republican from Lubbock). It left House and Senate committees without opposition, passed in the state House and Senate by wide margins in May, and was signed into law by the governor that June. As will become clear below, the bill's swift passage, occurring at the end of the school year, is helpful for our empirical strategy.

The bill had two primary features affecting current workers. First, it expanded the basis for computing retirement annuities from the highest three years of salary to the highest five. In effect, this reduces the pension annuity by the size of an experience step for a typical worker, usually a permanent three-percent reduction. For workers serving stints in leadership, or other highly paid roles, the cut was larger. For instance, a teacher serving three years as principal would see her annuity fall by 11 percent under the new regime.[5,6] Second, the reform enlarged the penalties for retiring early. Annuities for those claiming at age 60 fell by a third. Annuities for those claiming at age 55 were cut in half.[7]

The bill's sponsors argued that these cuts were a necessary response to ballooning liabilities and rising costs for schools and taxpayers. Rather than cuts, opponents of the bill proposed counter legislation raising annuities for current retirees and increasing the contribution of the legislature to the teacher pension fund by 42 percent, from 6 percent of payroll to 8.5 percent (House Research Organization 2005). The legislature passed the cost-reducing measure rather than the cost-increasing one.

The reform was well-communicated to teachers. In the summer after the bill passed,

[5]Calculating FAS based on more years mitigates the cost of pension "spiking." Pension spiking is the practice of having a few years of very high salary to substantially alter a worker's pension annuity (Fitzpatrick, 2017). The cannonical example is for workers to strategically time leadership to maximize annuities. Incorporating more years to compute a worker's final average salary (FAS) makes spiking harder and less effective at raising annuities above typical compensation.

[6]Teachers who served short stints in leadership experienced larger annuity cuts. Leadership pay in our data is typically an increase of 30 percent *ceteris paribus*. Because the final average salary was calculated on a larger base, FAS falls more for employees who had short stints in leadership. For example, a teacher who served as principal for the last three years of her career would see her yearly FAS fall by by 11 percent under the new regime, reducing her pension annuity in proportion.

[7]Teachers could retire with 20 years of service at age 55 in the old regime and receive 90 percent of their full benefits. After the reform, the same person would receive 47 percent of their full benefits.

the teacher union's magazine *Advocate* featured the pension reform as its cover story with a picture of a rotting apple core. In her quarterly letter to members, the union president wrote: "*Sine die*...it feels just that way—that someone or something has *died*. Could it be our hopes and dreams...? This was the most disappointing and disheartening legislative session we have ever seen" (emphasis original, Haschke (2005)).[8] To its hundreds of thousands of members (including 291,000 unionized teachers in Texas), the union provided detailed, accurate information about the reform and the grandfathering cutoff (Texas State Teachers' Association, 2005; National Center for Education Statistics, 2008).

To soften the blow, the reform grandfathered workers nearer retirement, leaving their pension benefits unaffected. A worker would be grandfathered if, on Sept 1, 2005, she were at least 50 years old, had at least 25 years of experience, or had age and service summing to at least 70. The relevant threshold for just over half of workers was the age cutoff. These are workers who began working for Texas public schools sometime after their 30th birthday. The relevant threshold for the substantial remainder was the rule-of-70 which grandfathered workers who started careers in public schools relatively early in life.[9] Employees who hadn't met one of those thresholds were entered in the new regime while those meeting at least one criteria were shielded from any change in their pension benefits.

To assess the size of the cut, we compare pension benefits under the two regimes. We calculate final average salary (FAS) for each worker in the administrative records we collected using a three- and five-year basis. On average, the FAS is 4 percent higher using three years than five, implying a 4 percent cut to retirement annuities. Workers see larger cuts when they have a few years of a high-salary position since a larger base reduces the the influence of irregular high years.

To examine the effect of greater penalties for early retirement on pension benefits, we calculate the present value of retirement wealth under both regimes for various archetypal workers.[10] We calculate the present-value of benefits assuming workers (i) live to age 84 (the

[8] *Sine die* has the same orthography as the English verb "to die," but the two are unrelated. The Latin phrase refers to adjourning a proceeding with no appointed date to resume. It literally means "without a date".

[9] The experience cutoff was not relevant for most workers because a worker would have had to begin continuous employment before age 20 for experience to be the relevant cutoff.

[10] We hold constant the years of service to show the effect of the reform on pension wealth.

mean for college-educated women), (ii) have an annual discount factor of 4 percent (Ericson and Laibson, 2018; Johnston, 2021), and (iii) have final average salary of $94,500 in 2021 dollars (the sample mean in 2005). For workers with 10 years of experience at age 50, the annuity reduction comes only from the expansion of the base for calculating final average salary, usually reducing the annuity by 4 percent. The early-retirement cut never affects this group because they will be full-retirement eligible before becoming early-retirement eligible. For workers with 20 years of experience, a worker maximizes the present value of her pension wealth by claiming at age 55, generating a stream of benefits valued at $1.2 million, approximately $700,000 in present discounted value at age 55. After the reform, the penalties for early retirement cut the annuity stream in half if the teacher claims at age 55. The worker maximizes the present value of her retirement by delaying retirement for five years, which reduces the present value of her retirement stream by 18 percent compared to the old regime. (Assuming longer life expectancy or lower discounting generates a smaller perceived reduction.)[11] The cut is smaller for teachers who become early-retirement eligible between ages 55 and 60. For workers who are already eligible for full retirement at age 50, the cut is 4 percent from the expanded basis of the FAS calculation.

In the study of pensions, examining the employees of public schools presents several advantages. First, the majority of public-sector workers are in education, providing large administrative records and accompanying statistical power. Second, because of the profession's size, reach, and influence on long-run outcomes, the results are valuable in themselves (Chetty et al., 2014; Papay and Kraft, 2015). Third, data on private-sector workers are quite hard to collect. Since teachers and other workers in public schools are government employees, their employment records are often available to researchers. Fourth and finally, output measures are available in education in the form of granular achievement data, allowing analysts to examine how pensions affect output and performance.

[11] To demonstrate the size of the cut, these calculations fix years of service and FAS to compare the two regimes with the same inputs. If we allow years of service to respond endogenously, the present value of retirement income under the reform is actually 2 percent greater than that in the status quo, but it comes at the cost of working an additional five years. When we allow the FAS to evolve also the new regime's present value is 6 percent greater than the old.

3 Data Description and Background

We collect administrative staffing records on all public-school employees from the Texas Education Agency covering the state from 2000 through 2021. The data comprise the yearly employment records for 2,046,975 individual workers. Many are teachers (50 percent), but the data also include nurses, therapists, librarians, bus drivers, custodial staff, and other pension-eligible employees. In each year, the records indicate a worker's unique identifier, district, campus, age (measured on September 1 of each year), professional role, grade assignment, subject taught, years of experience, base pay, total pay, full-time or part-time status, and whether the worker is a contract worker.[12] The state provided a file on the education attainment and graduating institution of all staff. Yearly employment records allow us to observe when a worker leaves a school, a district, or public education in Texas altogether.

To examine output, we collect achievement data in math and English language arts (ELA) publicly reported by the Texas Education Agency at the school-grade level, for grades three to eight, from 2002 to 2011. This provides us with data three years prior to and seven years after the reform. We cannot link students directly to their assigned teachers within a school-grade cell, but we do know which teachers provide instruction in math and ELA and the grades to which they are assigned.

In Texas, workers are eligible for normal retirement if they are 65 years old with at least 5 years of service, or the sum of their age and experience total 80 or more (e.g., a 50 year old worker with 30 years of experience is eligible for retirement).[13] When she meets the requirement, the worker is eligible to claim a pension annuity, which is calculated:

$$A_i = YOS_i \times 2.3\% \times FAS_i \tag{1}$$

The annuity received by an eligible worker i is a product of the worker's years of service

[12]We do not have workers' exact dates of birth which would suggest an alternative approach using the discontinuity in age, as measured in days around the 50-year threshold and the cutoff date of September 1. The state is concerned about safeguarding the privacy of employees.

[13]Since 2005, the state pension has undergone further reform, but these reforms have not affected the cohorts we study since they grandfather essentially all incumbent workers.

(YOS_i), the state's benefit multiplier (2.3% in Texas), and the average salary of a worker's highest earning three (or five) years, imprecisely called "final average salary" (FAS_i). A worker retiring this year with 30 years of credit and a final average salary of $80,100 (the sample mean) would receive a yearly annuity of $55,269 for the rest of her life, adjusted for inflation each year. A worker can claim a reduced annuity if she is at least 55 years old with at least 5 years of experience. For each year before normal retirement, a penalty is applied to the worker's annuity. Before the reform, early-retirement penalties were approximately 2 percent for every credit short of full-retirement eligibility (if she retired at age 55 with 20 years of service, she would be 5 years short of normal eligibility and the annuity would be reduced by 5 $\times$ 2% = 10% of the annuity calculated in equation (1)). In the new regime, the penalties were raised to 6–8 percent per year, so the penalty was multiplied by a factor of three or four.[14]

4 Regression Discontinuity and Worker Retention

The primary purpose of the state pension is to retain workers in the state's public schools (Morrissey 2017; Weller 2017). The panel nature of the data allows us to observe when a worker departs public-school teaching in Texas, precisely the outcome of interest. If a worker moves out of state or to a private school, she has not been retained in the public system and disappears from our records. As our main measures of retention R_i^y, we observe whether an individual i is still working in public schools y years after the reform.

Using the age and experience a worker has in the fall of 2005, we calculate each worker's distance to the grandfathering cutoff. To do so, we calculate three values: (1) the employee's distance beyond the age cutoff (age on September 1 2005 minus 50), (2) the employee's distance beyond the experience cutoff (experience accrued by September 2005 minus 25), and (3) the employee's distance beyond the rule-of-70 cutoff (experience plus age in September 2005 minus 70). The worker need only meet one of these notches to be grandfathered, so a worker's effective distance to grandfathering is the most positive

[14]The penalties aren't strictly formulaic and are circulated in tables to public sector workers (Teacher Retirement System of Texas, 2019).

distance to any notch. Those with distance greater or equal to zero are grandfathered and those with negative values are subject to the reform.

We model the outcome variable R_i (usually retention for y years or average retention odds) as a continuous function of distance to the grandfathering cutoff, and we estimate the outcome discontinuity that occurs at the threshold:

$$R_i = \beta T_i + f(x_i - x') + u_i \tag{2}$$

Here, $x_i - x'$ is each worker's distance to the eligibility cutoff, x', and T_i equals one if worker i was subject to the reform and zero if she wasn't. Thus, $f(x_i - x^{'})$ is a function of the running variable that captures the continuous relationship between distance to the cutoff and departure, allowing for unbiased estimation of the effect of the discrete policy change. We first collapse the data to the distance-to-the-cutoff level, which is asymptotically equivalent to clustering the standard errors on the running variable (Lee and Card, 2008; Johnston and Mas, 2018).

We present both linear and quadratic specifications to accommodate the possibility of curvature in the outcomes. The estimates from the two models are never (statistically) distinguishable, but the point estimates are smaller when using a quadratic specification. We calculate Imbens-Kalyanaraman bandwidths and use a triangular kernel, as recommended (Imbens and Kalyanaraman, 2012). Because the quadratic specification requires additional degrees of freedom, we sometimes have too few observations to estimate standard errors with the optimal bandwidth. Therefore, we use twice the IK bandwidth when estimating quadratic models and show estimates with a wide range of alternative bandwidths to assess robustness. Each alternative bandwidth produces very similar estimates throughout the support when estimating quadratic models, and we find estimates are quite stable within a sizeable range for linear models. At larger bandwidths, estimates for linear models tend to produce very large estimates which we view as specification error.

4.1 Evaluating the Regression Discontinuity

The regression discontinuity produces an unbiased estimate if two assumptions hold. First, the assignment variable (distance to the threshold) must have a continuous effect on the outcome at the cutoff. This isn't directly testable, but it is supported by the fact that outcomes are generally smooth away from the cutoff along the assignment variable. The second assumption is that no unobserved determinants of the outcome are discontinuous at the cutoff. In examining education and pension legislation in Texas, we find no other laws or reforms that rely on the same or similar cutoff rules (i.e., age 50, 25 years of service, or the rule of 70) to determine benefits during work or in retirement.

The primary threat to identification is if workers can somehow manipulate their placement around the threshold with false age or experience values, introducing a discontinuity in unobservables driven by selection. We believe this is far-fetched, as it would be committing fraud, and age and experience are easily verified in records already available to the state. There is also no empirical evidence to suggest this behavior occurred. If there were an excess (deficit) mass of workers on the favorable (unfavorable) side of the threshold, it would suggest that some workers were able to manipulate their assignment variable to avoid the reform (McCrary, 2008).[15] We present the density of workers around the cutoff in Figure 1. The distribution is smooth, including at the threshold, suggesting there was no scope for selection by manipulation of the running variable.

4.2 Retention Estimates from the Discontinuity

Figure 2 presents worker retention as a function of a worker's distance to the threshold. We display four series here, each indicating the share of workers that have been retained over time after the reform. At the cutoff, there is little visible discontinuity in retention rates. We present formal RDD results in Table 2. We estimate the effect of the reform on various measures of retention: whether a worker is retained for at least 1 year, at least 5 years, at least 10 years, and at least 15 years (the longest reach allowed by the time that has

[15] Alternatively, the Texas legislature could have chosen the exact cutoffs to benefit particular subpopulations of education workers. The fact that they simply chose round numbers (50, 25, 70) suggests this was not the case.

passed since the reform at this writing). To generate an single measure of retention odds, we calculate the *mean* retention rate over 15 years for each worker we observe in the reform year.[16] The discontinuity estimates reflect the effect of the reform for workers at the cutoff. The point estimates are usually positive, meaning the reform appears to increase retention. The estimates are significant in the linear model, but not quadratic.

Local linear estimates suggest that the reform increased five-year retention by 1.9 percentage points (2.3 percent) and ten-year retention by 1.9 percentage points (3.6 percent), both significant at the 0.001 level. In Appendix Figure A.3 we show estimated retention effects each year after the reform. Effects are close to zero in first few years, after which the reform increases retention and the effect remains at about 2 percentage points from five to eleven years post reform—years in which early retirement was most penalized. The average worker at the cutoff is 49.4 years old in the year of the reform, so five to eleven years after the reform reflects the period when reform was affecting workers concentrated around ages 54.4 to 60.4, the ages at which early retirement was attractive in the old regime but unattractive in the new one. After eleven years, the effects fall, eventually becoming insignificant, but they remain positive through the end of the observation window. Over the 15-year window, the reform is associated with an average 1.2 pp increase in retention (2.2 percent, significant at the 0.001 level). In the quadratic specification, the point estimate on the effect of the reform on five- and ten-year retention suggests the reform increased retention by 1.2–1.6 percentage points with 95-percent confidence intervals ruling out small attrition effects (ruling out larger than zero percentage points for five-year retention and ruling out larger than 0.02 percentage points for ten-year).

We present estimates at a range of bandwidths (from two to twenty, where the baseline bandwidth is usually 3 for the optimal bandwidth and 6–7 for twice optimal) to explore how bandwidth selection affects the estimates, which we present in Appendix Figures A.4–A.7. The linear estimates are robust to reductions or expansions of the bandwidth. The optimal bandwidth tends to be conservative, with larger estimates at both narrower and broader bandwidths. All estimates are positive and highly significant. The estimates are stable as

[16] Specifically, we take the mean of fifteen variables indicating whether a worker has been retained in each year.

the bandwidth varies in quadratic models. Each of the estimates is small, positive, and about half are insignificant at the five percent level. To assess whether unobservables differ across the threshold, we estimate placebos that apply the reform cutoff to workers in 2000, before the reform was implemented. We difference the placebo estimates from those in the treatment year, presented in Appendix Table 1. Four of five baseline estimates are robust (1-year retention, 10-year retention, 15-year retention, and average retention rate), but the five-year retention effect is not. Exploration reveals that the five-year retention discontinuity appears to be the product of a natural kink in five-year exit rates (because five years after the reform many treated teachers are then eligible for full-retirement), and this kink presents a specification challenge.[17] The overall estimate suggests that the reform increased retention by 1-2 percentage points after differencing discontinuities, similar to the baseline estimate.

5 Examining the Impact on Worker Output

We next turn to measuring the effect of the reform on worker output. Our identification strategy is based on quasi-experimental variation in the percentage of teachers affected by the reform across school-grade cells. We construct a measure for what fraction of teachers in each school-grade cell were exposed to the reform in 2005. The treatment variable indicates what share of the teachers in that grade were exposed to the reform in fall 2005 as a result of the age and experience of each teacher in the cell. If all the teachers met (missed) the grandfathering cutoff, the treatment variable is zero (one) for that school-grade cell. The average cell has 70 percent of its teachers affected by the reform, with a standard deviation of 35 percent.

We link these school-grade cells to the yearly achievement measures of the same unit. To estimate the effect of the reform, we implement a generalized difference-in-differences design:

[17]The reason for the kink is intuitive. Individuals to the right of the threshold often had met the rule of 70, those on the left had not. In five years all those that had met the rule of 70 will now satisfy the rule of 80 if they've worked continuously, and so are eligible for full retirement. Thus the right side of the cutoff are full-retirement eligible in five years, which generates a kink.

$$A_{sgt} = \alpha_s + \gamma_g + \delta_t + \beta(\mathbf{1}(t \geq 2005) \times T_{sg}) + \varepsilon_{sgt} \quad (3)$$

Here, A_{sgt} represents the achievement of students in school s in grade g at time t. α_s reflects a vector of school-specific fixed effects, γ_g captures grade-specific differences, and δ_t are year fixed effect which control for secular trends. $T_{sg} \in [0, 1]$ measures the treatment exposure of grade g in school s at the time of the reform. The parameter of interest is the coefficient on $(\mathbf{1}(t \geq 2005) \times T_{sg})$ where $\mathbf{1}(\cdot)$ is a function indicating that time is after the reform. Implicitly, we control for stable unobservable differences across treatment groups while adjusting for changes in common unobservables using the control group (cells that were completely shielded from the reform). The coefficient β reflects how achievement changes after the reform in treatment cells compared to control cells. The outcome reflects the fraction of students who met the state proficiency standard. We also estimate saturated models that include campus-year fixed effects. In these specifications, the identifying variation comes from within-school comparisons across grades with different exposures to the reform. Because the treatment-share of a cell can change over time, we instrument the actual treatment share with the treatment share predicted using the composition of teachers in the year before the reform.

In Table 3, we present the results of the difference-in-differences estimation. In the basic specifications, the reform is associated with achievement gains. 2SLS yields an estimate suggesting that treatment increased pass rates by 1.0 percentage point in math (where the average passage rate is 78.5 percent, effect significant at the 0.1 percent level) and 0.4 percentage points higher in reading (where the average passage rate is 80.8 percent; not significant). When controlling for school-year effects, or accounting for a cohort's average proficiency rate from the previous year, we find null effects of the reform on student achievement. We perform a version that exploits treatment close to the threshold which yields very similar estimates.[18]

[18] We generate a variable indicating what fraction of a school-grade cell are treated within a ten-year radius of the cutoff and condition on the fraction of a school-grade cell that is treated or untreated within a ten-year radius. These estimates are also small and insignificant (not presented).

6 Conclusion

In this study, we use variation in the pension regime workers face to examine the effect of a pension cut on retention and output. For grandfathered workers, pension benefits carry on as before. For workers under age and experience cutoffs, pension annuities are lower in retirement and penalties for retiring early are much higher.

The empirical design presents a few analytic benefits. First, the variation in regimes that we study is credibly exogenous. Rather than endogenously selected by workers, the policy altered the regime of workers based on attributes they could not control without foreknowledge. Second, the setting provides a tidy control group of like workers who were not affected by the reform. Implicitly, the estimates compare similar workers in a shared setting, where the only difference is their pension scheme. Third, the public-school setting provides data on retention using administrative employment records. Fourth, in most settings we can't hope to gauge worker productivity. In schools, records on output are readily available.

The reform has small positive effects on worker retention, probably because the reform makes early retirement less attractive. A body of work suggests that pay cuts result in lower morale and productivity. Another suggests the cuts implemented would elicit greater effort by raising the stakes of dismissal. Using measures on worker output, we find that the reform is not associated with declines in achievement. The results represent a rare opportunity to examine similar workers working in the same setting who yet were part of different pension regimes.

References

Anzia, S. F. (2017). "Pensions in the Trenches: Are Rising City Pension Costs Crowding Out Public Services?". *Working Paper*, pages 1–50.

Aubry, J.-P. and Crawford, C. V. (2017). State and local pension reform since the financial crisis. *State and Local Pension Plans*, (54).

Barro, R. J. (1991). Economic growth in a cross section of countries. *Quarterly Journal of Economics*, 106(2):407–443.

Biasi, B. (2019). "Higher Salaries or Higher Pensions? Inferring Preferences from Teachers' Retirement Behavior". *Working Paper*.

Brainard, K. and Brown, A. (2018). Significant reforms to state retirement systems. *National Association of State Retirement Administrators*, pages 1–98.

Brown, K. (2013). "The Link Between Pensions and Retirement Timing: Lessons from California Teachers". *Journal of Public Economics*, 98:1–14.

Chetty, R., Hendren, N., Kline, P., and Saez, E. (2014). "Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States". *The Quarterly Journal of Economics*, 129(4):1553–1623.

Costrell, R. M. and McGee, J. B. (2010). Teacher pension incentives, retirement behavior, and potential for reform in arkansas. *Education Finance and Policy*, 5(4):492–518.

Czernich, N., Falck, O., Krestchmer, T., and Woessmann, L. (2011). "Broadband Infrastructure and Economic Growth". *The Economic Journal*, 212(552):505–532.

Duflo, E. (2001). "Schooling and the Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment". *American Economic Review*, 91:795–813.

Ericson, K. M. and Laibson, D. (2018). Intertemporal choice. *NBER Working Paper 25358*, pages 1–67.

Fitzpatrick, M. D. (2015). "How much are public school teachers willing to pay for their retirement benefits?". *American Economic Journal: Economic Policy*, 7(4):165–188.

Fitzpatrick, M. D. (2017). Pension-spiking, free-riding, and the effects of pension reform on teachers' earnings. *Journal of Public Economics*, 148:57–74.

Fitzpatrick, M. D. (2019). Pension reform and return-to-work policies. *Journal of Pension Economics Finance*, 18(4):500–514.

Friedberg, L. and Webb, A. (2005). "Retirement and the Evolution of Pension Structure". *Journal of Human Resources*, 40(2):281–308.

Fuchsman, D., McGee, J. B., and Zamarro, G. (2020). "Teachers' Willingness to Pay for Retirement Benefits: A National Stated Preferences Experiment.". *Working Paper*.

Goldhaber, D., Grout, C., and Holden, K. L. (2017). Pension structure and employee turnover: Evidence from a large public pension system. *ILR Review*, 70(4):976–1007.

Gustman, A. L., Mitchell, O. S., and Steinmeier, T. L. (1994). The role of pensions in the labor market: A survey of the literature. *ILR Review*, 47(3):417–438.

Hanif, F., Millard, C., Bass, E. J., Curmi, E., Lam, D., and Pittaway, N. (2016). "The Coming Pensions Crisis: Recommendations for Keeping the Global Pensions System Afloat". pages 1–127.

Haschke, D. N. (2005). "Reflections on the 79th Legislature". *Advocate*, (Summer):2.

Imbens, G. and Kalyanaraman, K. (2012). "Optimal Bandwidth Choice for the Regression Discontinuity Estimator". *The Review of Economic Studies*, 79(3):933–959.

Johnston, A. C. (2021). Preferences, selection, and the structure of teacher pay. *Working Paper*.

Johnston, A. C. and Mas, A. (2018). Potential unemployment insurance duration and labor supply: The individual and market-level response to a benefit cut. *Journal of Political Economy*, 126(6):2480–2522.

Koedel, C. (2019). "California's Pension Debt is Harming Teachers and Students Now–And It's Going to Get Worse".

Koedel, C. and Xiang, P. B. (2017). Pension enhancements and the retention of public employees. *ILR Review*, 70(2):519–551.

Kube, S., Marechal, M. A., and Puppe, C. (2013). "Do Wage Cuts Damage Work Morale?". *Journal of the European Economic Association*, 11(4):853–870.

Lazear, E. P. (1979). Why is there mandatory retirement? *Journal of Political Economy*, 87(6):1261–1284.

Lee, D. S. and Card, D. (2008). Regression discontinuity inference with specification error. *Journal of Econometrics*, 142(2):655–74.

Legislative Budget Board (2005). Fiscal note, 79th legislative regular session: Sb1691 by duncan. pages 1–4.

Manoli, D. and Weber, A. (2016). “Nonparametric Evidence on the Effects of Financial Incentives on Retirement Decisions”. *American Economic Journal: Economic Policy*, 8(4):160–182.

McCrary, J. (2008). “Manipulation of the Running Variable in the Regression Discontinuity Design”. *Journal of Econometrics*, 142(2):698–714.

Mitchell, O. S. (2020). “Building Better Retirement Systems in the Wake of the Global Pandemic”. *No. w27261 National Bureau of Economic Research.*

National Center for Education Statistics (2008). Total number of public school teachers and percentage of public school teachers in a union or employees' association, by state: 1999-2000, 2003-04, and 2007-08.

Ni, S. and Podgursky, M. (2016). How teachers respond to pension system incentives: New estimates and policy applications. *Journal of Labor Economics*, 34(4):1075–1104.

Ni, S., Podgursky, M., and Wang, X. (2020). Teacher pension plan incentives, retirement decisions, and workforce quality. *Journal of Human Resources*, pages 1218–1232.

Ni, S., Podgursky, M., and Wang, X. (2021). Teacher pension enhancements and staffing in an urban school district. *Working Paper 21-01.*

Papay, J. P. and Kraft, M. A. (2015). “Productivity Returns to Experience in the Teacher Labor Market: Methodological Challenges and New Evidence on Long-Term Career Improvement”. *Journal of Public Economics*, 130:105–119.

Quinby, L. D., Sanzenbacher, G. T., and Aubry, J.-P. (2018). How have pension cuts affected public sector competitiveness? *Center for Retirement Research*, (59).

Sabatosa, J. (2021). Teachers oppose pension plan. *Rutland Herald.*

Teacher Retirement System of Texas (2019). Percentages to be applied to amount of standard service retirement benefit.

Texas State Teachers' Association (2005). “Changes in Retirement”. *Advocate*, (Summer):10–11.

Weakliem, D. L. and Frenkel, S. J. (2006). “Morale and Workplace Performance”. *Work and Occupations*, 33(3):335–361.

World Economic Forum (2017). “We'll Live to 100–How Can We Afford It?”. *White Paper*, pages 1–24.

Figure 1: Density of the Assignment Variable around the Cutoff

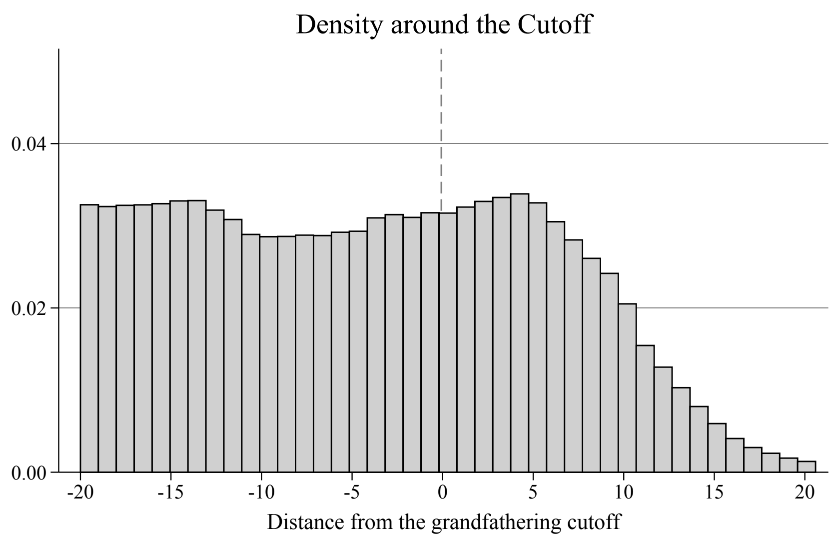

Note: Here we present the distribution of workers around the grandfathering cutoff for all workers employed in the year of the reform. An excess mass on the favorable side of the discontinuity would signal manipulation of the running variable. We detect no such manipulation. Source: Administrative data from Texas Education Agency.

Figure 2: Retention Patterns around the Cutoff

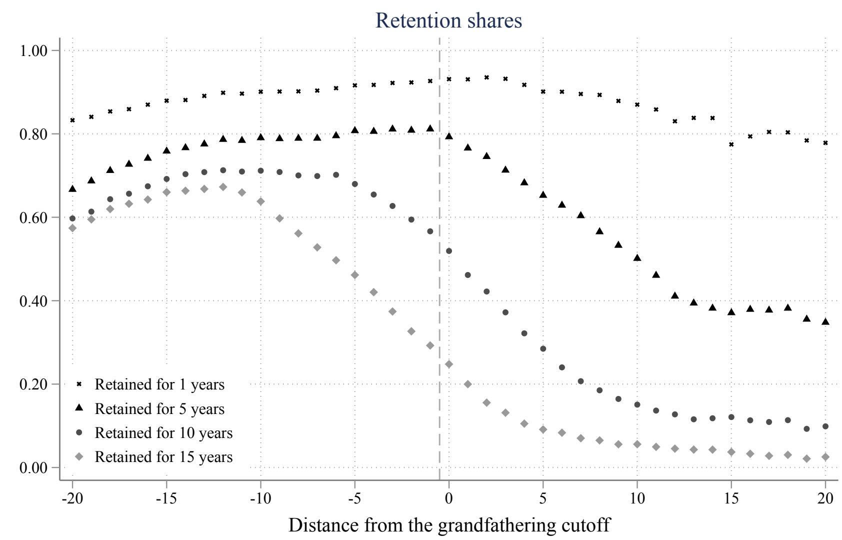

Note: In this figure, we present the share of workers who remain employed by public schools in Texas for $y \in [1, 5, 10, 15]$ years after the reform. Each dot represents the average retention rate for workers with a given distance to the grandfathering cutoff. Retention appears smooth across the cutoff. Source: Administrative data from Texas Education Agency.

Table 1: Summary Statistics

	(1) Mean	(2) SD	(3) N
Age	44.20	11.41	620,508
Experience	8.035	9.808	620,508
Here in 1yr	0.883	0.322	620,508
Here in 5yrs	0.696	0.460	620,508
Here in 10yrs	0.508	0.500	620,508
Here in 15yrs	0.377	0.485	620,508
Distance	-4.842	10.83	620,508
Grandfathered	0.373	0.484	620,508
Teacher	0.499	0.500	620,508
Base pay	$33,492	$16,876	620,508
Other pay	$949	$2,076	620,508
Bachelors	0.568	0.495	620,508
Masters	0.182	0.386	620,508

Notes: This table presents summary statistics for the analytic sample which includes all the workers we observe in the year of the reform.

Table 2: Regression Discontinuity Estimates of Retention Effects

	(1) Retained ≥1 year	(2) Retained ≥5 years	(3) Retained ≥10 years	(4) Retained ≥15 years	(5) Average Ret. Rate
Linear Controls					
Reform (RDD)	-0.001	0.019***	0.019***	0.007	0.0120***
	(0.002)	(0.002)	(0.004)	(0.006)	(0.0011)
Bandwidth	2.77	3.70	3.61	3.60	3.25
Quadratic Controls					
Reform (RDD)	0.000	0.016*	0.012^{+}	0.000	0.0063
	(0.006)	(0.005)	(0.006)	(0.007)	(0.0037)
Bandwidth	5.54	7.39	7.22	7.19	6.49
Mean DV (at cutoff)	0.929	0.809	0.529	0.248	0.549
Observations	620,508	620,508	620,508	620,508	620,508

Notes: This table presents the RDD estimates of the effect of pension reform on worker retention. We use the IK-optimal bandwidth and a triangular kernel in the linear specification, and twice the optimal bandwidth for the quadratic specification (independent quadratic terms on either side of the cutoff). We present the constant so the reader can gauge the size of each effect relative to the counterfactual. Other coefficients are omitted to spare clutter. + p<0.10, * p<0.05, ** p<0.01, *** p<0.001

Table 3: Difference-in-Differences Estimates of Output Effects

	(1) First stage	(2) Reduced form	(3) Two-stage least squares	(4)	(5)
	Reform	Achvmt	Achvmt	Achvmt	Achvmt
Mathematics					
Reform x post		0.333*	0.960***	0.265	0.261
		(0.133)	(0.331)	(0.352)	(0.427)
Instrument	0.919***				
	(0.006)				
N	50,734	50,734	50,734	36,890	34,057
R-squared	0.509	0.681	0.681	0.717	0.856
Reading					
Reform x post		0.280*	0.433	0.392	-0.210
		(0.140)	(0.352)	(0.385)	(0.447)
Instrument	0.921***				
	(0.007)				
N	49,152	49,152	49,152	36,871	34,295
R-squared	0.494	0.600	0.600	0.636	0.836
School FE	X	X	X	X	
Year FE	X	X	X	X	
Grade FE	X	X	X	X	X
Pre-scores				X	X
School-year FE					X

Notes: This table presents instrumental-variable difference-in-differences estimates of the effect of pension reform on worker output. Column (1) reports the first stage. Here, the outcome variable is the percent of reform-affected teachers in each school-grade in each year and the instrument is the percent of reform-affected teachers in the year before the reform is announced (2005). Column (2) reports the coefficient on treatment interacted with an indicator for post reform from equation 3, reflecting an estimate of the intent-to-treat effect of the reform. Treatment is the fraction of a grade that was treated in 2005. Columns (3) through (5) report estimates from a two-stage least squares DID. The estimated coefficient is that on a $reform \times post$ variable where $reform$ indicates the share of teachers in the grade who are subject to the reform. The instrument is the fraction who were treated in that cell before the reform was announced, interacted with post. The outcome variable in columns (2), (3), (4), and (5) is the fraction of students who meet the proficiency standard in math or reading. In math, the average is 78.5, and the standard deviation is 16.6. In reading, the average is 80.8, and the standard deviation is 14.8. School-year FE attenuate the estimates to insignificance, as do controls for the previous scores of each grade. + p<0.10, * p<0.05, ** p<0.01, *** p<0.001

A Additional Exhibits

Figure A.1: Non-Parametric Regression Discontinuity, Five-Year Retention

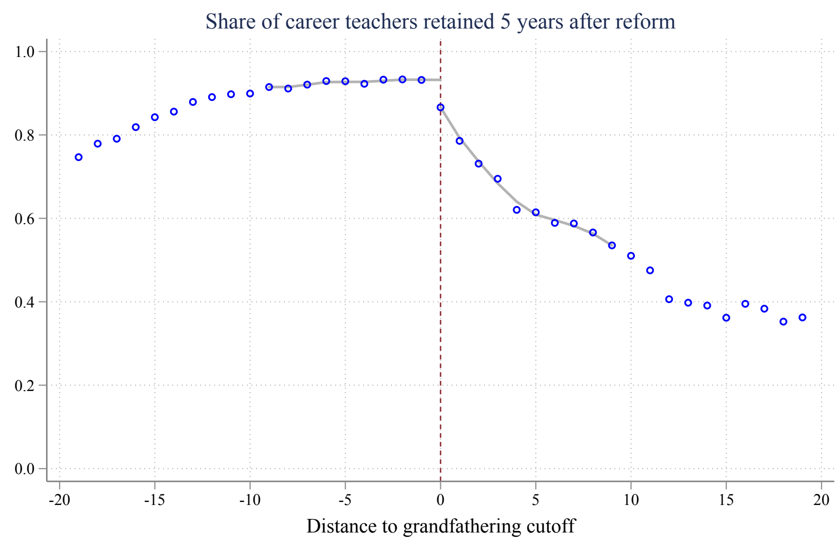

Note: In this figure, we show the raw data used to estimate the five-year retention effect of the reform. Source: Administrative data from Texas Education Agency.

Figure A.2: Non-Parametric Regression Discontinuity, Ten-Year Retention

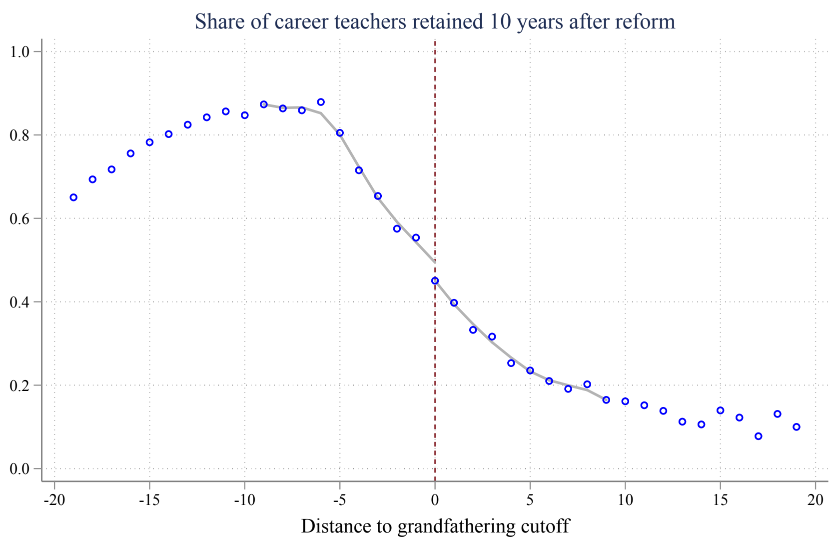

Note: In this figure, we show the raw data used to estimate the ten-year retention effect of the reform. Source: Administrative data from Texas Education Agency.

Figure A.3: Effect over Time

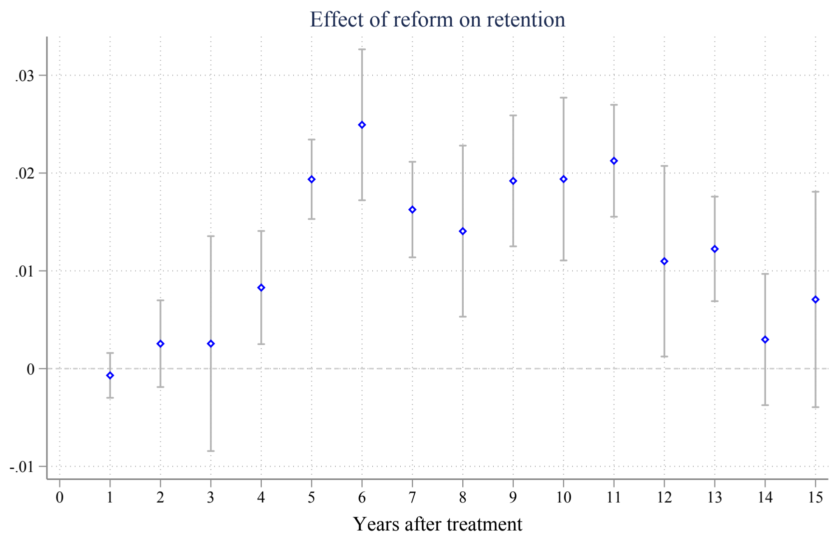

Note: In this figure, we estimate the effect of the reform on retention at various years after the reform. We estimate the RDD with a triangular kernel and an independent linear terms on either side of the threshold within the optimal bandwidth. The effects are indistinguishable from zero immediately after the reform, grow to approximately 2 percentage points from five years after the reform through eleven years after the reform, and shrink back to approximately zero. Source: Administrative data from Texas Education Agency.

Figure A.4: Assessing Bandwidth Selection

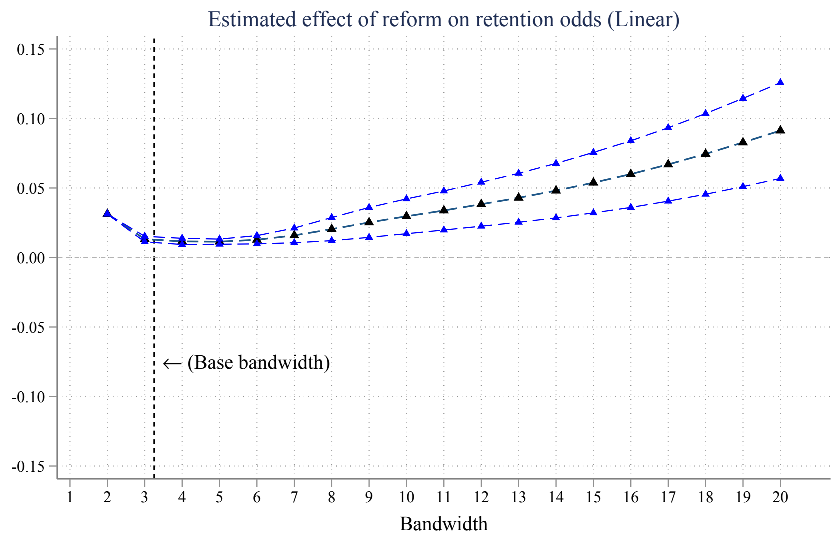

Note: In this figure, we estimate the regression discontinuity while varying the bandwidth to explore the sensitivity of the estimates to this choice. The outcome variable is the average retention rate (seen in Column (5) of Table 2). We estimate the RDD with a triangular kernel and an independent linear terms on either side of the threshold. From a bandwidth of 3 to 7, the estimates are quite consistent with our main estimate reported in Table 2. Larger retention estimates are retrieved when bandwidths are smaller or larger than this range. Source: Administrative data from Texas Education Agency.

Figure A.5: Assessing Bandwidth Selection

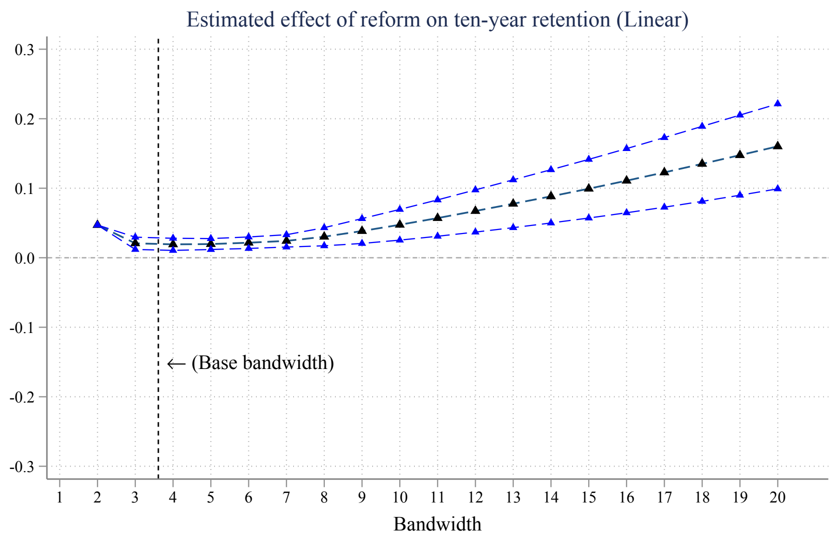

Note: In this figure, we estimate the regression discontinuity while varying the bandwidth to explore the sensitivity of the estimates to this choice. The outcome variable is the ten-year retention rate (seen in Column (3) of Table 2). We estimate the RDD with a triangular kernel and an independent linear terms on either side of the threshold. From a bandwidth of 3 to 8, the estimates are quite consistent with our main estimate reported in Table 2. Larger retention estimates are retrieved when bandwidths are smaller or larger than this range. Source: Administrative data from Texas Education Agency.

Figure A.6: Assessing Bandwidth Selection

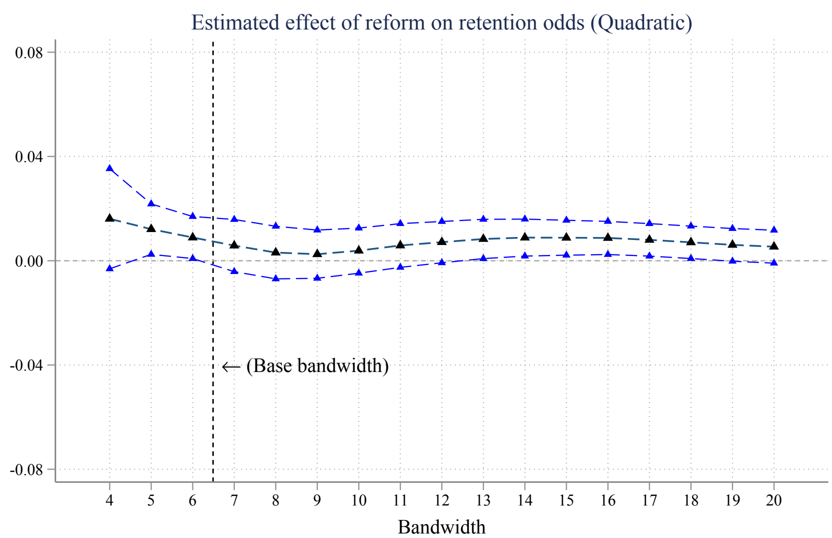

Note: In this figure, we estimate the regression discontinuity while varying the bandwidth to explore the sensitivity of the estimates to this choice. The outcome variable is average retention rates (seen in Column (5) of Table 2). We estimate the RDD with a triangular kernel and an independent quadratic polynomial on either side of the threshold. Throughout the bandwidth space, the estimates are quite consistent, always positive, and are statistically significant at the five-percent level in about half of estimates. Source: Administrative data from Texas Education Agency.

Figure A.7: Assessing Bandwidth Selection

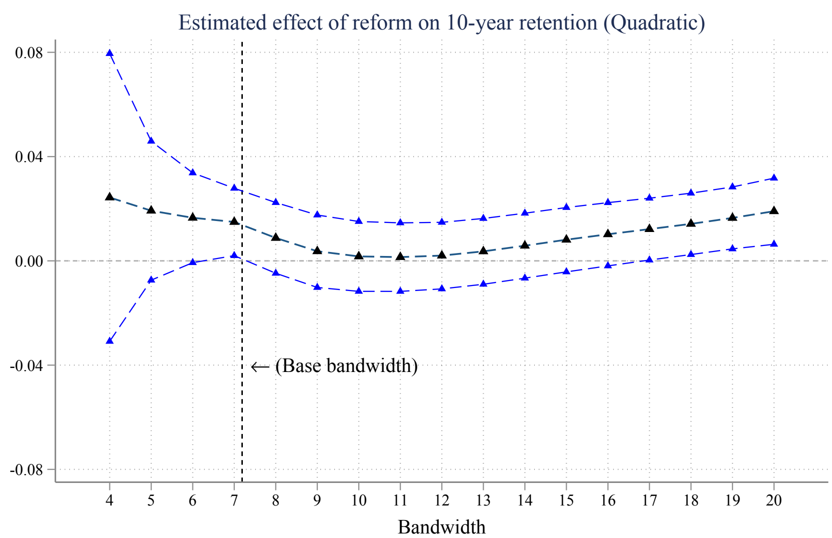

Note: In this figure, we estimate the regression discontinuity while varying the bandwidth to explore the sensitivity of the estimates to this choice. The outcome variable is average retention rates (seen in Column (3) of Table 2). We estimate the RDD with a triangular kernel and an independent quadratic polynomial on either side of the threshold. Throughout the bandwidth space, the estimates are quite consistent and always positive. Source: Administrative data from Texas Education Agency.

Figure A.8: Distribution of Treatment Exposure

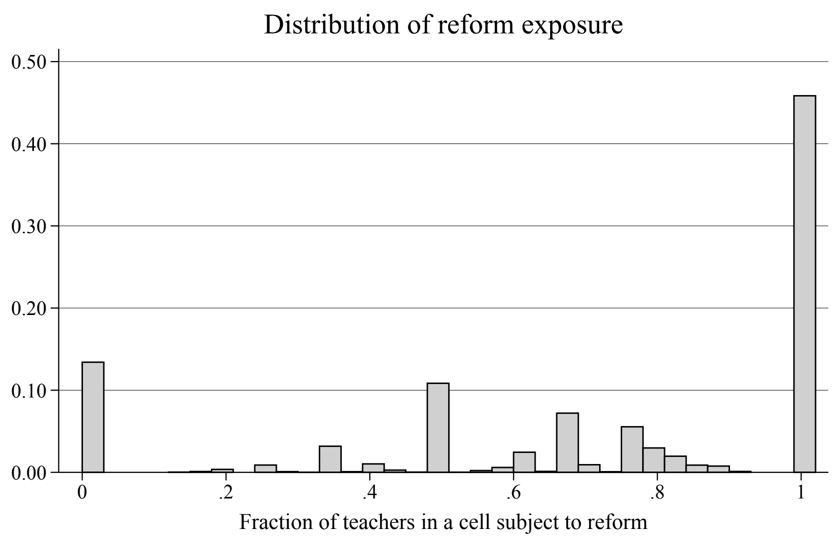

Note: In this figure, we present the distribution of the reform exposures of school-grade cells in 2005. On average, 70 percent of teachers in a cell are exposed, and exposure ranges from 0 to 100 percent. Source: Administrative data from Texas Education Agency.

Figure A.9: First Stage for Each Year

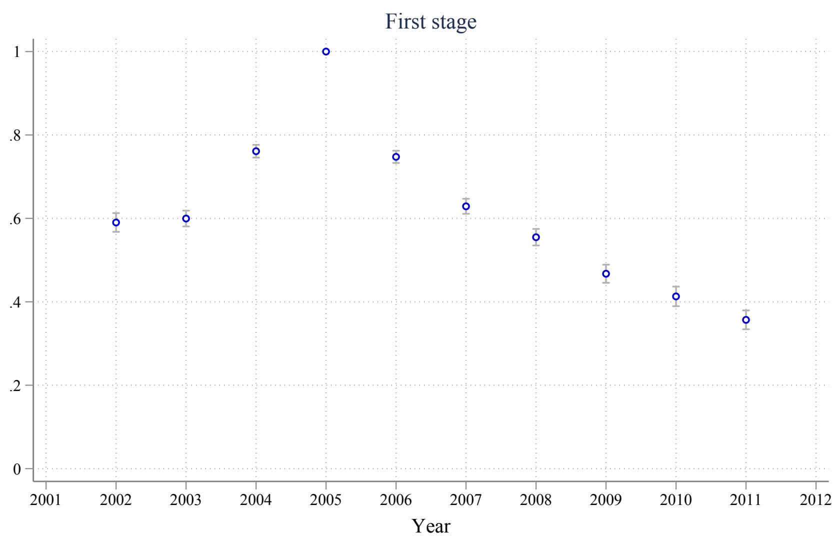

Note: In this figure, we plot the coefficient of the instrument (share of a school-grade's teacher who were subject to the reform in September 2005) in regressions of the share of school-grade teachers subject to the reform. Source: Administrative data from Texas Education Agency.

Figure A.10: Dynamic Difference-in-Differences with Two-Stage Least Squares

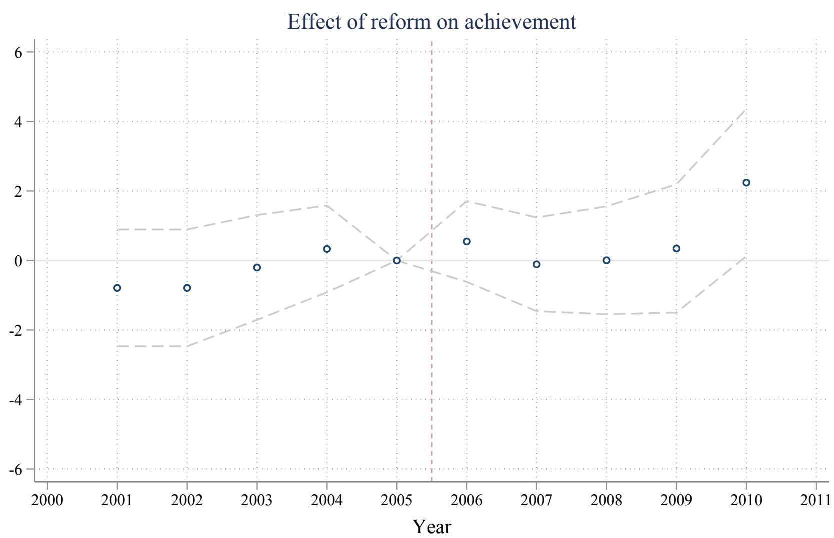

Note: In this figure, we plot the coefficient of predicted treatment shares for a cell based on the instrument for each year. Source: Administrative data from Texas Education Agency.

Table 1: Placebo RD and Difference-in-Discontinuities Estimates

	(1) Retained ≥1 year	(2) Retained ≥5 years	(3) Retained ≥10 years	(4) Retained ≥15 years	(5) Average Ret. Rate
Panel A: All workers (N=620,508)					
Treatment RD	-0.001	0.019***	0.019***	0.007	0.0120***
	(0.002)	(0.002)	(0.004)	(0.006)	(0.0011)
Placebo RD	-0.005*	0.021**	0.001	0.013**	0.0010
	(0.002)	(0.006)	(0.006)	(0.004)	(0.0050)
Difference-in-RD	0.005	-0.001	0.018**	-0.006	0.0110*
	(0.003)	(0.007)	(0.007)	(0.007)	(0.0051)
Panel B: Career workers (N=125,956)					
Treatment RD	0.003***	0.072***	0.058***	0.015	0.0316***
	(0.000)	(0.010)	(0.019)	(0.010)	(0.0061)
Placebo RD	0.007**	0.062***	-0.007*	-0.015*	0.0120***
	(0.002)	(0.014)	(0.003)	(0.006)	(0.0030)
Difference-in-RD	-0.004	0.010	0.065***	0.030*	0.0196*
	(0.002)	(0.017)	(0.019)	(0.012)	(0.0068)

Notes: This table presents RDD estimates of the effect of pension reform cutoff in the main sample and in a placebo sample using workers in 2000, five years before the policy. In the third row of each panel, we present the difference-in-RD estimates which is the first row subtracted by the second row. All estimates use a local-linear specification with a triangular kernel. + p<0.10, * p<0.05, ** p<0.01, *** p<0.001

Table 2: Regression Discontinuity Estimates of Retention Effects among Teachers

	(1) Retained ≥1 year	(2) Retained ≥5 years	(3) Retained ≥10 years	(4) Retained ≥15 years	(5) Average Ret. Rate
Linear Controls					
Reform (RDD)	-0.002***	0.029***	0.019**	0.004	0.0102***
	(0.000)	(0.005)	(0.006)	(0.004)	(0.0014)
Bandwidth	2.76	3.72	3.67	3.91	3.43
Quadratic Controls					
Reform (RDD)	0.001	0.024*	0.007	-0.002	0.0034
	(0.006)	(0.008)	(0.010)	(0.010)	(0.0060)
Bandwidth	5.52	7.44	7.33	7.83	6.86
Mean DV (at cutoff)	0.953	0.859	0.513	0.211	0.623
Observations	309,860	309,860	309,860	309,860	309,860

Notes: This table presents the RDD estimates of the effect of pension reform on worker retention among teachers in the data. We use the IK-optimal bandwidth and a triangular kernel to estimate the linear specification, and twice-optimal bandwidth with triangular kernel for the quadratic specification. We present the constant so the reader can gauge the size of each effect relative to the counterfactual. Column (1) suggests attrition effects in year 1 from the reform. This appears to be sensitive to local bandwidth choices. At any of the bandwidths used for the other coefficients, the coefficients are very near zero and statistically insignificant. Other coefficients are omitted to spare clutter. + $p<0.10$, * $p<0.05$, ** $p<0.01$, *** $p<0.001$