Journal of Pension Economics and Finance (2021), **20**, 151–168
doi:10.1017/S1474747219000362

CAMBRIDGE
UNIVERSITY PRESS

ARTICLE

The trade-off between pension costs and salary expenditures in the public sector

Dongwoo Kim[1], Cory Koedel[2]* and P. Brett Xiang[2]

[1]Economics, Texas Christian University, Fort Worth, Texas, USA and [2]Department of Economics and Truman School of Public Affairs, University of Missouri, Columbia, Missouri, USA
*Corresponding author. Email: koedelc@missouri.edu

(Received 28 February 2019; revised 22 October 2019; accepted 24 October 2019; first published online 10 January 2020)

Abstract

We examine pension-cost crowd out of salary expenditures in the public sector using a 15-year data panel of state teacher pension plans spanning the Great Recession. While there is no evidence of salary crowd out prior to the Great Recession, there is a shift in the post-recession years such that a 1% (of salaries) increase in the annual required pension contribution corresponds to a decrease in total teacher salary expenditures of 0.24%. The effect operates through changes to the size of the teaching workforce, not changes to teacher wages. An explanation for the effect heterogeneity pre- and post-recession is that public employers are less able to shield the workforce from pension costs during times of fiscal stress. This problem is exacerbated because unlike other benefit costs, such as for health care, pension costs are countercyclical.

Key words: Pension costs; pension crowd-out of salaries; pension incidence; teacher pensions

JEL Codes: H75; J32; I22

1. Introduction

The balance sheets of public defined-benefit (DB) pension plans have been steadily worsening since the turn of the century. Plans covering public educators are no exception. Using data from the Bureau of Labor Statistics' National Compensation Survey, Costrell (2015) estimates that inflation-adjusted per-pupil pension costs paid by employers have more than doubled from \$500 in 2004 to over \$1,000 by 2015, on average. The cost increase is driven predominantly by rising unfunded liabilities (Backes *et al*., 2016).

As a matter of accounting, pension costs are paid by public employees and their employers, but the true incidence is unclear. States and municipalities can respond to the rising costs by raising revenue or borrowing, which puts incidence on taxpayers (or future taxpayers), and/or by cutting expenditures in other areas of the budget, which reduces the quality and availability of services and can affect worker salaries and other benefits. There is little research aimed at understanding how rising pension costs are being managed by state and local governments, despite growing awareness of the significant impact of these costs on budgets (e.g., Belvedere, 2016; Tucker, 2017; Petrella and Pearson, 2019).

We contribute to the literature by estimating the incidence of pension costs on salary expenditures for covered public-sector employees. Although reducing salary expenditures for covered workers is just one way to offset rising pension costs, it is an intuitive one, and one that follows naturally from how pension plans collect contributions (i.e., on a 'per-head' basis among covered employees). Moreover, our focus on teachers is of interest because research suggests they do not value their DB pension benefits highly (Fitzpatrick, 2015), at least at the margin, yet teacher DB plans are consistently supported by labor groups (Weingarten, 2017).[1] It may be that the incidence of pension costs on teacher salary expenditures is small, which would help to explain the support DB plans receive from organized labor.[2] But if the incidence is high, it would raise questions as to why teacher labor groups continue to support public DB pension plans.[3]

We examine pension-cost incidence using a 15-year national data panel of state teacher plans. Our data on pension-plan finances are taken primarily from plans' actuarial valuation reports (AVRs) and comprehensive annual financial reports (CAFRs), which we supplement in some cases with data from the public plans database (PPD) (maintained by the Center on Retirement Research at Boston College). We merge these data with state data on teacher salaries and employment levels from the National Center for Education Statistics (NCES). Our data panel covers the years 2001–2015, spanning the Great Recession.

Our empirical analysis is closely related to parallel literature on the incidence of health benefit costs (Clemens and Cutler, 2014; Lubotsky and Olson, 2015). We estimate models that link within-state fluctuations in pension contribution costs to fluctuations in salary expenditures. We also estimate models that allow for a change in the relationship between pension costs and salary expenditures after the 2008 financial crisis. To reduce concerns about endogeneity, we favor models that leverage variation in plans' actuarially calculated annual required contributions (ARCs) for identification, rather than contributions that are actually paid. The latter can be manipulated politically. In principle, our use of the ARC should allow for clean identification, although we note potential limitations below (which we explore empirically to the extent possible).

Our aggregate models that combine pre- and post-recession data yield suggestive evidence of pension crowd out and our confidence intervals are in line with what has been found for crowd-out in the recent health-benefit literature, but none of the estimates is statistically significant. However, when we allow for differential crowd-out in the pre- and post-recession periods, our models reveal important heterogeneity. Whereas prior to the Great Recession there is no evidence that pension-cost fluctuations influenced teacher salary expenditures, in the post-recession period a one-percentage-point increase in the ARC (i.e., 1% of teacher salaries) corresponds to a 0.24% reduction in total salary expenditures. Supplementary models show that the negative relationship between pension costs and salary expenditures is driven by changes to the size of the teaching workforce (relative to the counterfactual) and not changes to average teacher salaries. This implies that the quality of educational services has been reduced and teacher workloads have gone up as a result of rising pension costs since 2008.

An explanation for the pre- and post-recession heterogeneity is that during times of fiscal stress, government agencies are less able to tap revenues and resources from other parts of the budget in order to protect the workforce from rising benefit costs, including pension costs. An aspect of pension financing that likely contributes to this problem is the countercyclical nature of pension cost increases (Yin and Boyd, 2018). Specifically, when a recession hits the value of assets in the pension plan falls. This lowers the asset-to-liability ratio and triggers higher contributions at a time when public

[1]Fitzpatrick estimates that senior teachers are willing to give up very little in current salary to increase pension compensation – 20 cents per present-value dollar of pension benefits – at least at the margin. Studies by Chingos and West (2015), Clark et al. (2016), and Goldhaber and Grout (2016) find that teachers are split in terms of their preferences for defined-contribution (DC) versus DB retirement benefits, with a large fraction of teachers typically opting into each type of plan when given the choice.

[2]Glaeser and Ponzetto (2014) develop a political-economy model consistent with there being low incidence of pension costs on public workers. Indeed, there is evidence from the literature on employee health benefit costs showing that salary incidence is far less than one (Baicker and Chandra, 2006; Clemens and Cutler, 2014; Lubotsky and Olson, 2015; Anand, 2017), and pension costs are easier to 'shroud' than health costs because health costs are not deferred.

[3]Even if the incidence on teachers is high and teachers on average do not value their pensions at the cost, the positive position of labor groups on pensions could reflect the underrepresentation of young workers in these groups (Koedel et al., 2013). Unions also leverage pension plan assets to promote other union goals (Even and Macpherson, 2014) and teacher labor groups may not know the incidence of pension costs on salary expenditures (we are not aware of any previous evidence linking pension costs to salary expenditures in the literature).

employers are already facing tight constraints. This aspect of pension financing makes it distinct from the financing of other employee benefits, most notably health benefits, which do not depend on investment returns and are thus less responsive to macroeconomic conditions.

2. Background

2.1 General information

Most public-sector employees, including public school teachers, receive retirement benefits in the form of a DB pension, paid as a lifetime annuity. Pension payments are a function of the final average salary (typically calculated as the average salary over the highest few years of earnings in covered employment) and years of plan-covered service. Service years are multiplied by a 'formula factor,' which is often around 2% for teachers, to yield a replacement rate of the final average salary in retirement. For example, a 30-year worker with a 2-% formula factor would receive 60% of her/his final average salary annually. Plans vary in how cost-of-living adjustments are incorporated after retirement.[4]

Public pension plans are designed to be funded at the cohort level. In principle, the idea is that each year contributions paid by employees, and contributions paid on their behalf by employers and states, collectively cover the actuarial present value of benefits accrued for that year. The cost of providing retirement benefits for services performed by current workers is referred to as the 'normal cost' by pension actuaries.[5]

The ARC to fund a pension plan, calculated as a percentage of worker salaries, is a combination of the normal cost and the cost of paying down the unfunded actuarially accrued liability (UAAL). The UAAL is effectively debt from past plan operations. It is amortized as dictated by the accounting rules of a plan, which set the amortization window and determine whether the UAAL is re-amortized each year.

Actual contributions to the plan on behalf of teachers are made by a combination of the teachers themselves, their employers (school districts), and state governments. The contributing groups vary across states, as do their formal contribution shares. Actual contributions are typically linked to the ARC by state statute, sometimes with constraints (e.g., a statute might limit how fast the contribution rate can rise annually regardless of how the ARC is changing), but this is not always the case. As a counterexample, in California prior to 2014, the contribution rate was set by state statute and there was no mechanism by which it could increase or decrease in response to changes in the ARC.

2.2 Why are pension contribution rates rising across the USA

The ARC is the product of actuarial calculations that make many assumptions along dimensions including, among other things, life expectancies, career longevities, and career wage profiles. The most controversial assumption is the high assumed rate of return on assets, which is also the rate used to discount liabilities. In recent history, this rate has typically been around 8% in public plans nationally. While some states have reduced the nominal rate in recent years; on average across the USA, the *real* assumed rate of return has gone up because of downward adjustments to inflation assumptions (Biggs, 2018). Because worker benefits are guaranteed, financial economists have argued that a risk-free real rate is more appropriate (Novy-Marx and Rauh, 2009, 2011, 2014; Biggs, 2011; Munnell *et al.*, 2015).

The high assumed rate of return is made more problematic by asymmetric responses to periods of below- and above-average performance of the investment portfolio. Koedel *et al.* (2014) show that teacher plans across the USA implemented retroactive, unfunded benefit improvements in the late

[4]Some teachers are also enrolled in Social Security. Nationally, estimates of the share of teachers enrolled in Social Security (in addition to their state retirement plans) range from 60 to 75% (Kan and Aldeman, 2014; Koedel and Podgursky, 2016).

[5]Cohort-level funding still permits significant resource transfers between workers within a cohort, such as between short-career and long-career workers (e.g., see Friedberg and Webb, 2005; Costrell and Podgursky, 2010; McGee and Winters, 2013).

1990s and early 2000s on the heels of an impressive stock-market boom. The boom temporarily inflated pension fund balance sheets and the benefit improvements were rationalized by plans' better-than-average investment returns. Compensating benefit cuts did not occur after the 2001 stock market correction. Moreover, even pension reductions that have occurred since 2008 – when they have occurred – are smaller in magnitude than the benefit improvements documented by Koedel *et al.* (2014). Responding asymmetrically to periods of above- and below-average returns in this way ensures long-term debt accrual even if the pension portfolio makes the high assumed rate of return on average.

Because of funds' persistent inability to meet the assumed rate of return on assets, made worse by asymmetric modifications to benefit formulas, UAALs have been accruing rapidly in most state and municipal plans across the USA. As liabilities rise, the asset-to-liability ratio in a plan falls (all else equal), which triggers increases in the ARC. Backes *et al.* (2016) document that for new teachers as of 2015, the average per-worker cost of servicing the unfunded liability reported by state plans in the USA was just over 10% of salaries.

Table 1 shows the evolution of the ARC on average across state teacher plans during our data panel from 2001 to 2015.[6] Data from all plans dating back to 2001 are unavailable and thus the panel is unbalanced, but by 2003 most (41) states are included. The table documents a clear increasing trend in the ARC. Between 2003 and 2015, among a fairly stable sample of states, the ARC rose by over 50% on average, representing a substantial increase in labor costs. The variance of the ARC across states also increased over time as shown in column 3.

The rising trend in the ARC is driven predominantly by the cost of servicing UAALs. Figure 1 uses the Illinois plan as an illustrative example. The UAAL is fairly flat through 2007, then rises with the onset of the Great Recession and continues to grow thereafter. The ARC trend follows the UAAL trend closely. Note that in the aggregate data in Table 1, the average ARC begins to rise in the early 2000s, before the Great Recession. Although this could be partly explained by modest fluctuations in the economy, such as the 'dotcom crash' in the early 2000s, this explanation is unsatisfying because on the whole, from the mid-1990s (after the savings and loan crisis) through 2007 was a period of strong returns in financial markets and general economic prosperity in the USA. A better explanation is the combination of the high assumed rate of return and the widespread, unfunded benefit improvements in the late 1990s and early 2000s.

While changes to normal costs also impact the ARC trend in Table 1, the study reported that the plans' normal costs did not change quickly over the time period. Moreover, to the extent that they did change, they declined modestly. Factors influencing the decline are changes to actuarial assumptions and reductions to benefits for new entrants.[7] The decline due to benefit reductions is limited because benefits have not been reduced in all states. Moreover, where benefits have been reduced, incumbent teachers – like other public-sector employees – are grandfathered into the benefit structure into which they are hired and not affected.

3. Data and descriptive statistics

We construct a state-level data panel with information about the size of each state's teaching workforce, teacher salaries, and details about the DB pension plan covering public educators. Our panel

[6]The table includes the pension plan for teachers in Washington DC and excludes the pension plan in Alaska, which is the only plan in the nation without a DB component. Note the trend using a weighted average (weighted by active members) is very similar (not shown); the trend of the median ARC (not shown) is substantively similar but exhibits less growth (the median ARC grew by 63% from 2001–2015, compared to 82% for the average ARC), which reflects the fact that some states with high ARC values in the upper tail contribute significantly to growth in the average ARC.

[7]As an example of an assumption change that lowered normal cost, at one point Missouri plan reduced its inflation assumption but did not adjust the nominal expected rate of return on assets, which effectively increased the assumed real rate of return. Plans have made assumption changes that increase normal costs as well, such as direct reductions to the assumed rate of return on assets or updates of life tables that allow for greater longevity.

https://doi.org/10.1017/S1474747219000362 Published online by Cambridge University Press

Table 1. Average annual required contribution (ARC) across plans by year

Year	Obs.	Mean	Std. Dev.	Min.	Max.
2001	14	12.13	4.47	3.43	21.01
2002	31	13.39	5.54	3.36	25.74
2003	41	13.85	6.31	3.36	26.53
2004	45	14.89	6.07	4.9	28.45
2005	47	16.19	6.31	7.37	30.21
2006	47	16.84	6.03	7.7	29.44
2007	48	17.40	5.74	7.6	29.69
2008	49	17.48	5.90	8.32	30.71
2009	49	17.74	6.24	8.0	32.61
2010	50	18.68	6.79	7.5	33.62
2011	50	19.84	7.03	8.0	35.83
2012	50	20.69	7.22	8.0	35.85
2013	50	21.60	7.68	9.6	44.88
2014	50	22.04	7.34	9.84	45.39
2015	50	22.09	7.14	9.5	42.98

Notes: We use 49 states and Washington DC to populate this table. The table excludes the pension plan in Alaska, which is the only plan in the nation without a defined-benefit component. The years in the table are fiscal years; e.g., the year 2003 is the fiscal year that started on July 1, 2002 and ended on June 30, 2003.

includes 49 states and the District of Columbia.[8] Among the 50 plans, 27 cover educators exclusively, another 4 – Colorado, Kansas, Maryland, and Tennessee – have a formal educator division within a larger consolidated plan, and the remaining 19 states have consolidated plans that cover educators along with other public employees.[9] For separate educator plans and consolidated plans with a formal educator division, we use financial data relevant to educators. For the consolidated plans without a formal educator division, we use financial data for the whole plan. These decisions allow us to connect the relevant financial information about the pension plan for teachers to the labor outcomes we consider.

The financial information about the plans is taken primarily from plan-provided AVRs and CAFRs. In cases where the AVR and CAFR are not available, we supplement our dataset with the Public Plans Database (2001–2015) maintained by the Center on Retirement Research at Boston College.[10] Our data on total teacher salary expenditures, average teacher salaries, and the number of teachers employed in each state come from the NCES and cover public elementary and secondary schools.

Table 2 provides descriptive statistics on teacher salary expenditures, teacher employment, and basic economic conditions of states, overall and split into the pre- and post-recession periods. Total salaries and the average size of the teacher workforce decreased from the pre- to post-recession periods, and the cross-state variance in both variables also declined. Average salaries remained essentially flat. Economic conditions are as expected – the level and variance of unemployment across states increased in the post-2008 period and average state-level GDP is only modestly higher.

[8]Alaska is excluded because its DB pension plan was closed in 2006. For states with both DB and DC plans, we use data from the DB plan in our analysis. For states with both DB and hybrid plans, we combine the DB and the DB component of the hybrid plans. For states with only hybrid plans, we take the DB component of the hybrid plans. None of our findings is qualitatively influenced by reasonable adjustments to these procedures, including dropping states with hybrid plans or DC plan options.

[9]Non-certified educational staff (e.g., secretaries and maintenance) may or may not be covered by the same pension plans as teachers. For this analysis, we take information from the teacher plan whenever it is distinguishable in AVRs and CAFRs.

[10]The data taken directly from plan reports are preferable for two reasons. First, the PPD combines funding statistics in plans that cover multiple member groups, which prevents us from isolating the teacher potion of plans in some states. Second, we identify some discrepancies between the PPD and the data provided by plan reports, particularly in the early years of the data panel.

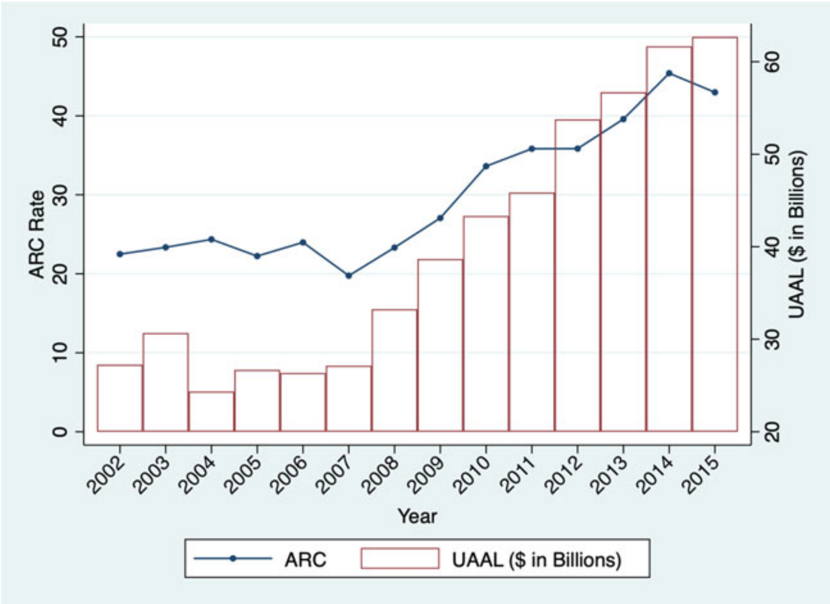

Figure 1. Trends in the annual required contribution (ARC) and unfunded actuarial accrued liability (UAAL) in Illinois.
Note: These data are as reported by the Illinois Teacher Retirement System.

4. Methodology

4.1 Core specifications

To begin, consider the following regression of teacher salary expenditures on the pension contribution rate (PCR) based on conceptually similar investigations of healthcare cost incidence by Clemens and Cutler (2014) and Lubotsky and Olson (2015):

$$\ln\left[Y_{j,t}(1-c_{j,t})\right] = \boldsymbol{X_{j,t}}\gamma_1 + \mathrm{PCR}_{j,t-k}\gamma_2 + \varphi_j + \theta_t + \omega_{j,t} \tag{1}$$

In Equation (1), $Y_{j,t}$ is total salary expenditures for state j in year t and $c_{j,t}$ is the teacher contribution rate to the pension plan, $0 \leqslant c < 1$. Thus, the dependent variable is the natural log of total salary expenditures net of teachers' own required pension contributions – loosely speaking, teachers' total 'take-home pay.' Netting out direct employee pension costs from the dependent variable is conceptually important because it allows the model to pick up cost pass-through via teachers' own contributions. For example, suppose that when total pension contributions increase, employee contributions are raised fully to cover the higher cost. If our dependent variable did not net out teachers' own pension contributions (i.e., if we replaced $Y_{j,t}(1-c_{j,t})$ with $Y_{j,t}$ in Equation (1)), the model would not detect the incidence on teacher salaries because NCES data measure salaries prior to teachers' pension contributions (like with other benefit payments teachers make). This is similar to how Lubotsky and Olson (2015) account for employees' own contributions to their health insurance premiums.[11]

The vector $\boldsymbol{X_{j,t}}$ includes time-varying state economic controls – namely, the log of GDP and the unemployment rate.[12] $\mathrm{PCR}_{j,t-k}$ is the variable of interest – the amount that was actually contributed

[11]While it is conceptually important to remove teacher contributions from the dependent variable in Equation (1) and subsequent models, substantively our findings are similar even if we do not. The implication is that teachers' own contribution rates are not immediately responsive to changes in pension costs at a scale that matters in our data (although this does not rule out a high degree of responsiveness in some states).

[12]The GDP data are from the Bureau of Economic Analysis and the unemployment data are from the Bureau of Labor Statistics.

Table 2. Descriptive statistics. Dollar figures are in real 2015 dollars

Variables	Full sample (2001–15) Mean	Std. Dev.	Pre-recession (2001–07) Mean	Std. Dev.	Post-recession (2008–15) Mean	Std. Dev.
Teacher salaries/employment						
Total salary (in millions)	4,723	5,698	5,261	6,283	4,354	5,233
Average salary	51,703	7,965	51,735	7,892	51,682	8,025
Total number of teachers	63,803	67,389	65,081	68,569	63,004	66,742
Economic characteristics						
Unemployment rate	6.05	2.03	4.89	1.04	6.86	2.15
GDP (in millions)	329,600	395,550	324,916	384,831	332,826	403,213
N	672		274		398	

Note: We use 49 states and Washington DC to populate this table. The table excludes the pension plan in Alaska, which is the only plan in the nation without a defined-benefit component. Teacher salary calculations are net of teachers' own pension contributions.

to the teacher pension plan in state j for year $t-k$ (we elaborate on timing issues below), as a percent of teacher salaries.[13] Because the total salary variable is in logs and the PCR is defined in units of salary percentage points, the estimate of γ_2 can be interpreted as an elasticity. For example, a point estimate of −0.01 for γ_2 would indicate that an increase in the PCR of one percentage point of teacher salaries corresponds to a 1% reduction in salary expenditures. φ_j and θ_t are state and year effects, respectively, which restrict the identifying variation to occur within states over time, conditional on the national time trend. The error term is ω in Equation (1), $\varepsilon_{j,t}$, is clustered at the state level (Bertrand *et al*., 2004).

Identification of γ_2 requires that variation in the PCR within states over time is exogenous. Some sources of variation in PCR seem promising in this regard; however, on the whole, interpreting the estimate of γ_2 from Equation (1) is complicated by potential endogeneity of the PCR. For example, consider a state in a financial crisis that actively chooses, among other things, to reduce its pension payments as part of its broader fiscal response. A choice under such circumstances would likely be correlated with state educational spending (along with another spending), which in turn could affect teacher salary expenditures.

Given the conceptual weakness of using the PCR directly in the model, we favor an approach linking teacher salary expenditures to variation in the ARC. Although the ARC may still be endogenous – an issue that we elaborate on in the next subsection over the course of discussing the identifying variation – it is not as susceptible to endogeneity as the PCR because it is purely an actuarial calculation. Equation (2) matches Equation (1), but replaces the PCR with the ARC:

$$\ln\left[Y_{j,t}(1-c_{j,t})\right] = X_{j,t}\beta_1 + A_{j,t-k}\beta_2 + \lambda_j + \tau_t + \varepsilon_{j,t} \tag{2}$$

This specification gives the reduced-form conditional relationship between the ARC and teacher salary expenditures.

Moreover, we can build on this approach by isolating ARC-driven variation in the PCR with instrumental variables analog to Equation (1), using the following two-stage model:

$$\mathrm{PCR}_{j,t-k} = X_{j,t}\alpha_1 + A_{j,t-k}\alpha_2 + \eta_j + \phi_t + u_{j,t} \tag{3}$$

$$\ln\left[Y_{j,t}(1-c_{j,t})\right] = X_{j,t}\pi_1 + \hat{\mathrm{PCR}}_{j,t-k}\pi_2 + \varsigma_j + \psi_t + e_{j,t} \tag{4}$$

[13]To construct the PCR variable we multiply the plan-reported ARC by the percent of the required contribution paid, as reported in the PPD. In many states the PCR is synonymous with the statutory contribution rate (SCR). We use the broader term PCR to reflect the fact that in the case of a discrepancy between the SCR and what was actually contributed, we capture the latter.

Equation (4) matches Equation (1) except that the parameter of interest, π_2, is identified using only ARC-driven variation in the PCR (through Equation (3)). Both the ARC and PCR are measured in percentage points of teacher salaries and consistent with intuition, the results from Equation (3) indicate a strong mapping between them. Specifically, we estimate the first stage coefficient α_2 with various values of k to be roughly 0.75. This implies that estimates of β_2 and π_2, from Equations (2) and (4) respectively, will be similar, which they are (see below).

Finally, we note that all of our models are estimated using state-level data because teacher pension plans are administered at the state level.[14] Although salary decisions are primarily local (i.e., made at the district level), using district-level salary data offers little value over our approach – which is effective to bring the outcome data to the same level of aggregation as the treatment data – given the state-level treatments and corresponding clustering structure.[15] This aspect of our study differs from that of Clemens and Cutler (2014) and Lubotsky and Olson (2015), who use district-level data in their analyses of healthcare cost incidence because health benefits are determined at the district rather than the state level.

4.2 Identifying variation

Because of the aforementioned endogeneity concerns with the PCR, our preferred models rely on variation in the ARC for identification. The ARC is the PCR actuaries calculates as necessary to pay for benefits accrued during the year and to service liabilities. It is an appealing measure of pension costs because it is based on an objective function of the fiscal health of a plan and plan accounting rules.[16] But while the ARC is an improvement over the PCR in the sense that the potential for endogeneity is less severe, endogeneity concerns with the ARC remain.

The potential for bias due to endogeneity of the ARC in our models must come from state-level dynamic factors that influence both the ARC and teacher salaries, which would be unaccounted for by the state fixed effects. Perhaps the most obvious issue is that pension plans overweight companies with in-state headquarters in their portfolios (Brown *et al.*, 2015). This is likely to induce a dynamic correlation between asset returns in a plan (and thus the ARC) and the local economy, which in turn could affect salary expenditures on teachers.

The time-varying economic controls in the X-vector are included in an effort to reduce bias from this source. In addition, below we show that our findings are substantively similar regardless of whether we include or omit these variables in our models, which indicates a limited potential for bias from omitted state-level economic controls. Upon further examination, the reason for the insensitivity of our estimates is that changes in economic conditions within states over time are only weakly correlated with changes in the ARC. To give an empirical sense of the relationship, we separately regress ARC_{jt} and GDP_{jt} on state and year effects to produce the residualized values $\mathrm{A\tilde{R}C}_{jt}$ and $\mathrm{G\tilde{D}P}_{jt}$, respectively, and then correlate the two series. The correlation is small and statistically insignificant (−0.026).[17] The weak correlation is not entirely unexpected because the own-state bias in

[14]There are five municipalities (Chicago, Kansas City, New York City, St. Paul, and St. Louis) across the country where teachers are enrolled in a separate plan from the larger state plan, and thus subject to different contributions, but these instances account for such a small fraction of teachers that they are ignorable. For completeness, we have confirmed that our findings are qualitatively robust to dropping the four states in which these five municipal plans reside. Moreover, as would be predicted if our measures of pension costs across the full teaching workforce are less precise in these states, our cost-incidence estimates disattenuate marginally (and insignificantly) when we drop them from the sample.

[15]This inference follows directly from Bertrand et al. (2004).

[16]Certainly in terms of intent this is an accurate description of the ARC, although we cannot rule out nefarious behavior. For example, pension boards could put pressure on actuaries, for whom the pension plans are clients, to set assumptions in order to produce a desired outcome (e.g., a better funded ratio). Such behavior would be hard to detect in pension plan documents but is possible.

[17]We test the significance of the correlation by regressing $A\tilde{R}C_{jt}$ on $G\tilde{D}P_{jt}$ with state clustering. The correlation is statistically insignificant (p-value=0.85). There is a similarly weak and insignificant correlation between the residualized state unemployment rate and the residualized ARC (correlation=0.157; p-value=0.11). State GDP fluctuations are a much stronger conditional predictor of teacher salary expenditures than fluctuations in the state unemployment rate (see below).

pension plan portfolios, while clearly present, is not overwhelming (Brown *et al*., 2015). Moreover, the ARC is typically calculated as a moving average, which smooths out fluctuations.[18]

Another concern with relying on the ARC for identification is that it can be influenced by past contributions and *future commitments* to contribute. Starting with the former, if the PCR is lower than expected by actuaries in year t, the ARC in year $t+1$ should rise, all else equal. A concerning scenario would be if the underpayment is endogenous – e.g., due to broader fiscal stress – per the discussion above, in which case a rising ARC after the period of fiscal stress would be correlated with (presumably) improving economic conditions. This, in turn, would induce positive bias in $\hat{\beta}_2$ and $\hat{\pi}_2$. However, as noted above, we see no evidence of a positive relationship between the residualized values of the ARC and the time-varying state economic variables, which lessens this concern. Moreover, the above-mentioned lack of sensitivity of our estimates to include the economic controls further supports the view that the scope for bias from this type of endogeneity is limited. Finally, we also note that any positive bias from this type of endogeneity should be strongest in the post-2008 period, but this is in the opposite direction of our findings. If anything, this suggests that the difference in the effect of pension costs on salary expenditures that we identify before and after the Great Recession is conservative.

Turning to the issue that ARC values in year t could be influenced by future commitments to contribute, a recent illustrative example is California. In 2014, California passed legislation to raise the PCR over a 7-year period. Because a (sizeable) component of the ARC is the amortized unfunded liability, future commitments to pay affect current actuarial calculations of the immediate ARC. This can lead to a situation in which a lower ARC today is the product of known increases in pension costs in the future. The concern in this scenario is that forward-looking school districts could react proactively, in which case we would see the ARC decline in year t, and in anticipation of a higher PCR in some year $t+k$, teacher salary expenditures would also decline in year t.[19]

While interesting conceptually, there does not seem to be much potential for this type of endogeneity in our application. We are not aware of any scheduled increases to the PCR like the California example prior to the 2008 financial crisis. Moreover, while many states have experienced substantial increases to the PCR since, we are not aware of other scheduled increases that are so explicit. In addition, as with the preceding issue, if bias from this source is present in our estimates, it is more likely to manifest in the post-recession years of our data panel given that PCR increases are more prevalent during that time.[20] This again suggests the potential for the post-recession estimates to be biased positively, which is in the opposite direction of what we find.

Stepping back from these specific endogeneity threats, it is useful to consider the factors that drive the within-state identifying variation in the ARC more broadly.[21] Factors that influence the ARC include investment returns on the pension portfolio, changes to actuarial assumptions, and changes to accounting rules. Demographic shifts could also influence the ARC, but only if they are sharp and unexpected. Predictable demographic shifts, such as a spike in retirements due to a bulge of retirement-aged workers, will not affect the ARC because such shifts are built into the actuarial calculations.

[18]Because of the smoothing, in results omitted for brevity we also estimated time-staggered correlations between residualized ARC and the economic variables that align previous-year economic conditions to contemporaneous ARC. We find similarly weak and insignificant correlations.

[19]Note that while the California example is useful for illustration, it does not impact our analysis in a meaningful way because the legislation was not introduced until 2014, leaving just 1 year of our data panel (2015) for which an anticipatory effect could occur.

[20]Many states have experienced sustained, rising PCRs that are primarily the product of state statutes tying the PCR to the ARC in some way. In these scenarios, school district administrators who understand pension accounting and their state statutes could predict future rising rates, but not precisely and there is no schedule of rising rates that could be consulted as in California.

[21]The within-state variation in the ARC we leverage for identification (conditional on the national time trend) accounts for about 15% of the total variance in the ARC in our data. We obtain the 15% number with a supplementary regression of the year-t ARC on state and year fixed effects, which yields an R^2 value of 0.85. The identifying variation we leverage proves sufficient to statistically detect elasticities in the range of 0.2-0.4 in absolute value, depending on the specification (see below).

On the whole, the above-mentioned sources of variation are generally appealing from an identification perspective. Perhaps the cleanest source of variation is year-to-year fluctuations in the investment return on the pension portfolio. In results that we omit for brevity, we attempted to reproduce our instrumental-variables models using the portfolio investment return as the instrument in Equation (3), but unfortunately, the first stage is too weak to have any traction. This implies that within-state variation in investment returns is insufficient to support our analysis alone. Changes to actuarial assumptions and accounting rules are also plausibly exogenous sources of variation in the ARC and these surely contribute to the variance we leverage in our models – e.g., updates to the life tables used by actuaries, or changes to the amortization period of the UAAL – but these changes are difficult to track comprehensively across plans and years so we were unable to develop an identification strategy that isolates specific assumptions or rule changes.[22] In sum, our estimates rely collectively on all of the variations in the ARC that occurs within states over time while conditioning on the national time trend.

4.3 Effect timing

With regard to timing, note that state plans report ARC values prospectively. For example, in Arizona, the AVR published in year $t-2$ reports the projected ARC for year t. We construct the data panel so that $k = 0$ indicates an ARC value that applies to school-year t, reported in year $t-x_s$, where x_s can vary by state. We report estimates from models where $k = 0, 1, 2$ because of uncertainty about the timing of any pension-cost effects. Forward-looking school districts with complete flexibility over wages and the employment level should respond to the year t ARC in year t, in which case $k = 0$ is appropriate. However, districts may not fully internalize projected cost changes until they hit the budget. This could result in effects that occur with a lag such that, for example, a value of $k = 1$ or $k = 2$ would capture the effect better.

A related timing issue is wage rigidity. Because labor negotiations occur at different times for different districts within a state, but pension coverage is centralized statewide, any effect of the ARC on salary expenditures using lag structure k, on average, will be the product of heterogeneous responses across districts within a state owing to differences in negotiating status. Put another way, in any year t, a fraction of districts will be negotiating teacher salaries and others will be locked into a contract. Conceptually, we would expect wage rigidity of this nature to attenuate our estimates and to the extent that this is an issue, a longer lag structure should allow us to observe larger effects. However, there is no evidence of larger crowd-out effects with a longer lag structure, which suggests that wage rigidity is unlikely to be suppressing our estimates. This interpretation is consistent with supplementary results below showing that the crowd-out mechanism changes in the size of the workforce, not average salaries.

4.4 Pre- and post-recession effect heterogeneity

Finally, we hypothesize that the incidence of pension costs may differ depending on macroeconomic conditions. To test this, we expand the models in Equations (2)–(4) to allow for effect heterogeneity in the pre- and post-recession years. The expanded version of Equation (2) is as follows:

$$\ln\left[Y_{j,t}(1-c_{j,t})\right] = X_{j,t}\delta_1 + A_{j,t-k}\delta_2 + (A_{j,t-k} \times \mathrm{POST}_t)\delta_3 + \xi_j + \kappa_t + \nu_{j,t} \tag{5}$$

Like terms in Equations (2) and (5) are as defined above and the variable POST_t is an indicator equal to one if the year is 2008 or later and zero otherwise. Note that the 'level effect' of POST_t is absorbed by the year fixed effects in the model (κ_t). Conditional on the level controls, the new

[22]State plans sometimes produce 'experience studies' that document changes to actuarial assumptions, but plan reports do not necessarily reference these studies when they occur and they are not always accessible, making it difficult to build a credible database of changes. Some changes also occur outside of experience studies and we are not aware of any comprehensive source documenting them.

Table 3. Results from reduced-form regressions of total salary expenditures, net of teachers' own contributions, on the (1) actual pension contribution rate (PCR) and (2) annual required contribution (ARC)

	(1)	(2)	(3)	(4)	(5)	(6)
Variables	Log (Total salary)					
PCR_t	−0.0002 (0.0003)					
PCR_{t-1}		−0.0002 (0.0003)				
PCR_{t-2}			−0.0001 (0.0003)			
ARC_t				−0.0006 (0.0015)		
ARC_{t-1}					−0.0013 (0.0016)	
ARC_{t-2}						−0.0019 (0.0021)
Unemployment rate	0.0075 (0.0057)	0.0080 (0.0057)	0.0071 (0.0057)	0.0076 (0.0057)	0.0084 (0.0057)	0.0074 (0.0055)
Log(GDP)	0.5131*** (0.1178)	0.5161*** (0.1206)	0.5172*** (0.1191)	0.5101*** (0.1155)	0.5126*** (0.1180)	0.5129*** (0.1158)
State FE	X	X	X	X	X	X
Year FE	X	X	X	X	X	X
R^2	0.9980	0.9981	0.9982	0.9980	0.9981	0.9982
Observations	672	622	572	672	622	572

Note: Standard errors are clustered at the state level.
***p < 0.01; **p < 0.05; *p < 0.1.

interaction term tests whether there is a change in the relationship between the ARC and teacher salary expenditures before and after the onset of the Great Recession.

We also estimate instrumental variables models analogous to Equation (5), following the structure of Equations (3) and (4). We use $A_{j,t-k}$ and $A_{j,t-k} \times \text{POST}_t$ as instruments for $\text{PCR}_{j,t-k}$ and $\text{PCR}_{j,t-k} \times \text{POST}_t$ in these models.

5. Results

5.1 Primary results

Table 3 shows estimated coefficients on the PCR and ARC variables from Equations (1) and (2). Each column is a separate regression. The first three columns show results from Equation (1) and the last three columns are for Equation (2). For each model, we show results using the three different lag structures ($k = 0, 1, 2$). The sample size declines as we increase the number of lags because doing so shrinks the effective size of the data panel.

The estimates in columns (4)–(6) using the ARC, which we prefer, are negative and of moderate magnitude, but are noisily estimated and not statistically significant. For example, taken at face value, our estimate when $k = 0$ implies that a one-percentage-point increase in the ARC corresponds to a 0.06% reduction in total salary expenditures. The upper-bound crowd-out estimate is around 35% given the size of the standard error. This estimate is similar to estimates from the health insurance literature based on similar models, in terms of both its magnitude and precision. For example, the analogous point estimate from the primary specification in Clemens and Cutler (2014) indicates a 15% reduction in salaries with a standard error of 33%.[23]

[23]Clemens and Cutler (2014) interpret their estimate as showing 'that roughly 15% of the cost of recent benefit growth was passed onto school district employees through reductions in wages and salaries' (abstract, pp.65). We interpret our findings less definitively owing to the imprecision of the estimates. Lubotsky and Olson (2015) find that a dollar increase in health benefit costs corresponds to a \$0.17 reduction in take-home compensation, with all of the adjustment coming through higher premium copayments. Baicker and Chandra (2006) differs from ours and the other studies along several dimensions, but they find no wage effect overall and a wage crowd-out effect of 23% among workers with employer-provided health insurance.

https://doi.org/10.1017/S1474747219000362 Published online by Cambridge University Press

https://doi.org/10.1017/S1474747219000362 Published online by Cambridge University Press

Table 4. 2SLS estimates of the effect of the actual pension contribution rate (PCR) on total salary expenditures net of teachers' own contributions, using the annual required contribution (ARC) as an instrument

	(1)	(2)	(3)
Variables	Log (Total salary)		
PCR_t	−0.0008		
	(0.0019)		
PCR_{t-1}		−0.0018	
		(0.0022)	
PCR_{t-2}			−0.0026
			(0.0028)
Unemployment rate	0.0078	0.0083	0.0063
	(0.0058)	(0.0056)	(0.0056)
Log(GDP)	0.5172***	0.5231***	0.5179***
	(0.1192)	(0.1202)	(0.1166)
State FE	X	X	X
Year FE	X	X	X
R^2	0.7524	0.7314	0.6939
Observations	672	622	572

Note: Standard errors are clustered at the state level.
***p < 0.01; **p < 0.05; * p < 0.1.

Table 5. Estimates of the effect of the annual required contribution (ARC) on total salary expenditures net of teachers' own contributions, allowing for effect heterogeneity in the pre- and post-recession periods

	(1)	(2)	(3)
Variables	Log (Total salary)		
ARC_t	0.0013		
	(0.0015)		
ARC_t × Post-recession	−0.0024*		
	(0.0013)		
ARC_{t-1}		0.0003	
		(0.0015)	
ARC_{t-1} × Post-recession		−0.0021	
		(0.0014)	
ARC_{t-2}			−0.0006
			(0.0019)
ARC_{t-2} × Post-recession			−0.0017
			(0.0015)
Unemployment rate	0.0079	0.0084	0.0070
	(0.0058)	(0.0057)	(0.0056)
Log(GDP)	0.5045***	0.5057***	0.5048***
	(0.1128)	(0.1157)	(0.1145)
State FE	X	X	X
Year FE	X	X	X
R^2	0.9980	0.9981	0.9982
Observations	672	622	572

Note: Post-recession variable is an indicator equal to one if the years is 2008 or later. Standard errors are clustered at the state level.
***p < 0.01; **p < 0.05; * p < 0.1.

Next, in Table 4 we show results from the instrumental variables models. The first stage output is reported in the appendix (Appendix Table A.1) – as noted above, the first stage shows a strong but imperfect correspondence between the PCR and the ARC. The first stage F-statistics are sufficiently large (ranging from 30 to 45 depending on the lag structure). The findings in Table 4 are similar to what we show in Table 3, although the estimates are somewhat more negative and less precise. None are approaching statistical significance at conventional levels.

Next, we allow for effect heterogeneity between the pre- and post-recession periods. First, in Table 5 we report on estimates from the model shown by Equation (5) – i.e., the reduced-form model where

Table 6. 2SLS estimates of the effect of the actual pension contribution rate (PCR) on total salary expenditures net of teachers' own contributions, allowing for effect heterogeneity in the pre- and post-recession periods, using the annual required contribution (ARC) and ARC × Post-recession as instruments

	(1)	(2)	(3)
Variables	Log (Total salary)		
PCR_t	0.0010		
	(0.0017)		
PCR_t × Post-recession	−0.0027*		
	(0.0014)		
PCR_{t-1}		−0.0003	
		(0.0019)	
PCR_{t-1} × Post-recession		−0.0022	
		(0.0016)	
PCR_{t-2}			−0.0015
			(0.0023)
PCR_{t-2} × Post-recession			−0.0018
			(0.0019)
Unemployment rate	0.0084	0.0082	0.0064
	(0.0060)	(0.0057)	(0.0055)
Log(GDP)	0.5162***	0.5135***	0.5059***
	(0.1175)	(0.1179)	(0.1143)
State FE	X	X	X
Year FE	X	X	X
R^2	0.7520	0.7352	0.7013
Observations	672	622	572

Note: Post-recession variable is an indicator equal to one if the years is 2008 or later. Standard errors are clustered at the state level. ***p < 0.01; **p < 0.05; *p < 0.1.

we insert the ARC directly as the independent variable of interest and include an interaction between the ARC and the post-2008 period. The pre/post-2008 split suggests a shift in the relationship between the ARC and total salary expenditures after the Great Recession. There is no evidence of a relationship in the pre-recession years for any lag structure. But we estimate a substantial, negative and significant (at the 10% level) coefficient on the ARC in the post-2008 years for $k = 0$, and substantively similar but statistically insignificant coefficients for $k = 1, 2$. The $k = 0$ coefficient indicates a differential crowd-out effect during the post-recession period of 24%.[24]

The results from the analogous instrumental variables regressions are substantively similar and reported in Table 6. Again, the output from the first-stage regressions is reported in the appendix (Appendix Tables A.2 and A.3). The post-period coefficients in these models range nominally from −0.0018 to −0.0027, corresponding to crowd-out effects of pension contributions on teacher salary expenditures of 18%–27%. Again only the $k = 0$ coefficient is statistically significant. As noted above, the negative relative coefficients for the post-2008 period in Tables 5 and 6 are obtained despite the increased potential for positive bias in these parameter estimates.

5.2 Robustness and extensions

We test the robustness of our findings to the inclusion of model covariates in Appendix Table A.4, which shows results from models that exclude the time-varying state economic characteristics in the X-vector. For brevity we show these results only for the reduced-form models using the ARC – i.e., from Equation (5) as reported in Table 5 – although the lack of sensitivity of our findings to including these controls holds up throughout our analysis. These results are consistent with the

[24]The total post-period crowd-out effects are mostly lower if we take the pre-2008 coefficients – which are statistically insignificant but nominally positive in two of the three scenarios – at face value. For example, in the $k = 0$ model, the total post-period crowd-out effect estimate is −0.0011, or (0.0013–0.0024).

Table 7. Estimates of the effect of the annual required contribution (ARC) on the average teacher salary net of own contributions, allowing for effect heterogeneity in the pre- and post-recession periods

	(1)	(2)	(3)
Variables	Log (Average take-home salary)		
ARC_t	−0.0012		
	(0.0013)		
ARC_t × Post-recession	0.0016		
	(0.0011)		
ARC_{t-1}		−0.0008	
		(0.0013)	
ARC_{t-1} × Post-recession		0.0014	
		(0.0011)	
ARC_{t-2}			−0.0004
			(0.0014)
ARC_{t-2} × Post-recession			0.0013
			(0.0012)
Unemployment rate	−0.0015	−0.0018	−0.0024
	(0.0036)	(0.0037)	(0.0035)
Log(GDP)	0.2186***	0.2106***	0.1962***
	(0.0548)	(0.0537)	(0.0472)
State FE	X	X	X
Year FE	X	X	X
R^2	0.9545	0.9562	0.9593
Observations	672	622	572

Note: Post-recession variable is an indicator equal to one if the years is 2008 or later. Standard errors are clustered at the state level. ***p < 0.01; **p < 0.05; *p < 0.1.

weak relationship between residual variation in the economic covariates and the ARC within states over time as documented above.

We also examine the robustness of our findings to restricting the analytic sample to the balanced panel of 41 state plans from 2003 to 2015 in Appendix Table A.5. The state and time fixed effects should in principle account for any composition effects associated with the unbalanced panel going back to 2001, but this robustness test is meant to provide additional assurance. We again highlight results from models that match those shown in Table 5. As anticipated, the results in Appendix Table A.5 are substantively similar to what we show in the main text, although the coefficients attenuate marginally and are no longer statistically significant.

Finally, we look for evidence that crowd-out is differentially present in consolidated versus teacher-only plans. As mentioned previously, in 27 states teachers are in their own plans, while in the other states teachers belong to larger consolidated plans that also cover other public-sector employees. It is reasonable to expect similar crowd out across different types of plans, although hypotheses that support differential effects can be imagined. For example, one possibility is that teachers are more politically powerful than other public-sector employees. In such a scenario, they may be better positioned to negotiate the effects of rising pension costs in plans exclusive to them. Appendix Table A.6 explores the potential for heterogeneity along this dimension, again replicating the basic structure of Table 5. The estimates lose precision when we split the data by plan type, but there is no indication that crowd-out rates differ between consolidated and teacher-only plans.

5.3 Mechanisms

We extend our analysis by examining two different ways that districts can distribute cost incidence within the 'salary expenditures' bundle: (1) reduce (or not increase) teacher salaries, and (2) reduce (or not increase) the number of teaching positions. In Tables 7 and 8 we replicate the analysis in Table 5, but divide total salary expenditures into these two parts: average salary and the total number of teachers (the total number of teachers is measured in full-time-equivalents as reported by the

Table 8. Estimates of the effect of the annual required contribution (ARC) on the total number of teachers, allowing for effect heterogeneity in the pre- and post-recession periods

	(1)	(2)	(3)
Variables	Log (Total number of teachers)		
ARC_t	−0.0002		
	(0.0015)		
ARC_t × Post-recession	−0.0029**		
	(0.0012)		
ARC_{t-1}		−0.0012	
		(0.0014)	
ARC_{t-1} × Post-recession		−0.0025**	
		(0.0012)	
ARC_{t-2}			−0.0002
			(0.0014)
ARC_{t-2} × Post-recession			−0.0020*
			(0.0011)
Unemployment rate	0.0004	0.0008	0.0002
	(0.0039)	(0.0041)	(0.0043)
Log(GDP)	0.2119***	0.2162***	0.2168***
	(0.0660)	(0.0629)	(0.0648)
State FE	X	X	X
Year FE	X	X	X
R^2	0.9978	0.9979	0.9981
Observations	672	622	572

Note: Post-recession variable is an indicator equal to one if the years is 2008 or later. Standard errors are clustered at the state level. ***p < 0.01; **p < 0.05; *p < 0.1.

NCES). This allows us to gain insight into the mechanism(s) by which total expenditures are reduced.

Table 7 shows results from models where we replace total salary expenditures with the state-average teacher salary in Equation (5). None of the coefficients on the ARC variables is statistically significant and there is no pre/post-recession shift. Moreover, the post-period coefficients are nominally positive. We find no evidence that average teacher salaries are influenced by ARC fluctuations in either the pre- or post-recession period.[25]

Next, we examine the effects on the size of the teaching workforce. Table 8 shows models analogous to Equation (5) but where the dependent variable is the number of PK-12 teachers, in logs. The table makes it clear that this is the margin for effect. While there is no significant relationship between the ARC and the size of the teaching workforce in the pre-recession period, in the post-recession period increases in the ARC correspond to a smaller workforce.[26]

The findings in Table 8 also raise questions about the interpretation of the null findings in Table 7 given the strong link between teacher experience and salaries. For example, if the workforce reductions in Table 8 are concentrated at certain points in the experience distribution (e.g., among inexperienced new hires), the experience profile of the workforce could change. In turn, and noting the reliance on experience in salary schedules across the USA, the composition effect would impact average salaries. Put another way, it is possible that the null findings in Table 7 mask a reduction in the average salary conditional on experience, hidden by a compositional shift toward more senior teachers in the workforce.

Unfortunately, the NCES data do not allow us to determine the extent to which the workforce reduction documented in Table 8 is from reduced hiring or increased exits. However, the Digest of Education Statistics (Snyder *et al.*, 2018) does provide intermittent salary data for teachers grouped

[25]It has been well documented that demographic changes among the teaching workforce are putting downward pressure on average wages nationally (Schmitz, 2016), but this general trend should be captured in the models shown in Table 7 by the year fixed effects.

[26]Lubotsky and Olson (2015) also find that school districts do not reduce the size of the workforce in response to rising health insurance costs during the pre-recession period (but their data do not go past 2008 to examine post-recession years).

by experience. These data are from the Schools and Staffing Survey, which was administered during the second half of our data panel in 2008 and 2012. We use these data to test whether the ARC and average teacher salaries – for fixed experience groups – are negatively related in the initial years of the post-recession period. If our null results from Table 7 are biased by compositional changes to the teaching workforce, we would expect to find evidence of average salary changes between 2008 and 2012 in response to the ARC holding experience constant. In results relegated to the appendix (Appendix Table A.7), we show that this is not what we find – there is no evidence that the ARC affects average salaries conditional on experience.[27] Although this evidence is not conclusive because the estimates are noisy (a problem made worse by the fact that only about two-thirds of states have experience-specific salary data), it is not consistent with the null findings for average salaries in Table 7 being driven by a workforce-reduction-induced composition effect.

6. Conclusion

Public DB pension costs are rising rapidly and all indications point toward continued increases due to unmet actuarial assumptions that result in the accumulation of unfunded liabilities (Novy-Marx and Rauh, 2009, 2011, 2014; Biggs, 2011; Munnell *et al.*, 2015). The incidence of rising pension costs is unclear and to the best of our knowledge, our study is the first to estimate cost incidence on salary expenditures for covered workers. Our empirical analysis is based on a data panel of 50 state plans that cover public educators over 15 years spanning the Great Recession.

We find no evidence of pension-cost incidence on salary expenditures prior to the Great Recession, but there is a shift in the post-recession period such that salary expenditures are reduced when pension costs rise. Supplementary models indicate that incidence is on the extensive rather than intensive margin – that is, the effect operates through changes in the size of the teaching workforce, not average teacher salaries. Thus, our findings imply that rising pension costs in recent years have resulted in larger workloads for teachers and reduced educational services, all else equal. It is not clear if the nature and level of incidence we document is enough to influence the historical support of DB pension plans by teacher labor groups (Weingarten, 2017). Fitzpatrick's (2015) low estimate of the marginal value of pension benefits for teachers (of 20 cents on the dollar) is similar in magnitude to the level of cost incidence we estimate since 2008, and our analysis suggests that salaries of incumbents (who are disproportionately likely to be represented by organized labor) are unaffected. Labor groups also have other reasons to support public pension plans (Even and Macpherson, 2014).

We also again note that our analysis focuses entirely on the incidence of pension costs with respect to expenditures on teacher salaries, but changes to salary expenditures represent one of many ways that states and school districts can respond to rising pension costs. Within the broader category of educational expenditures (and again noting that there is no reason that incidence must be constrained to occur within educational spending categories), a notable cost item worthy of consideration in future research is employee health benefits. Like salary expenditures, reducing health benefit costs in response to rising pension costs is intuitive. Health benefits are also controlled by individual districts, like salaries, and may be less rigid than salaries depending on the nature of labor negotiations.

The heterogeneity we document in pension-cost incidence between the pre- and post-recession periods can be explained by public agencies having limited options for dealing with rising pension costs during times of fiscal stress when less revenue is available. We also emphasize the countercyclical nature of pension-cost increases, which are triggered by declines in asset-to-liability ratios and thus sensitive to macroeconomic conditions. This likely exacerbates the recession-related effect heterogeneity we identify. It would be of interest to examine the uniqueness of this aspect of our findings to pension costs. Health benefit costs should be similarly difficult for government agencies to manage

[27]Specifically, Appendix Table A.7 shows results from a first-differenced regression of the log of average salaries for teachers with (a) 0–2 and (b) more than 20 years of experience, respectively. The first differencing is equivalent to including state fixed effects because we use just two time periods (2008 and 2012), and we also include the time-varying economic controls. The coefficients on the ARC are not significant for either experience group and are inconsistent in sign.

during recessions, but unlike with pensions, there is no reason to expect health benefit costs to fluctuate counter-cyclically.[28]

Finally, we conclude with a brief mention of another form of incidence that differs among teachers within the workforce and is not examined here. As noted previously, since the turn of the century changes to plan benefits have been in the direction of reduced benefits for new entrants. Incumbent teachers are grandfathered into the previous, more-generous benefit structures.[29] The monetary value of benefit reductions for new employees is not incorporated into our analysis, but the pension 'tiers' under which new hires have been enrolled in many state plans of late provide lower benefits to workers but require the same contributions (by employees and employers). Teachers covered by the more- and less-generous tiers work on the same salary schedules. Therefore, within the teaching workforce, younger workers have been disproportionately affected in terms of total compensation inclusive of the *value* of pension benefits.

Supplementary material. The supplementary material for this article can be found at https://doi.org/10.1017/S1474747219000362.

Acknowledgements. Kim is in the Department of Economics at Texas Christian University, Xiang is in the Department of Economics at the University of Missouri, Columbia, and Koedel is in the Department of Economics and Truman School of Public Affairs at the University of Missouri, Columbia. The authors thank Shawn Ni and Eric Parsons for useful comments. They gratefully acknowledge research support from the Laura and John Arnold Foundation. The views expressed here are those of the authors and should not be attributed to their institutions or funders. Any and all errors are attributable to the authors.

References

Anand P (2017) Health insurance costs and employee compensation: evidence from the national compensation survey. *Health Economics* **26**, 1601–1616.

Backes B, Goldhaber D, Grout C, Koedel C, Ni S, Podgursky M, Xiang PB and Xu Z (2016) Benefit or burden? On the intergenerational inequity of teacher pension plans. *Educational Researcher* **45**, 367–377.

Baicker K and Chandra A (2006) The labor market effects of rising health insurance premiums. *Journal of Labor Economics* **24**, 609–634.

Belvedere MJ (2016) Pension pipe dreams put taxpayers on hook: Analysis. *CNBC* (04.11.2016).

Bertrand M, Esther D and Sendhil M (2004) How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics* **119**, 249–275.

Biggs A (2011) An options pricing method for calculating the market price of public sector pension liabilities. *Public Budgeting and Finance* **31**, 94–118.

Biggs A (2018) Public Sector Pensions Assume Record-High Investment Returns. *Forbes* (12.04.2018). Retrieved on (02.18.2019) at: Available at https://www.forbes.com/sites/andrewbiggs/2018/12/04/public-sector-pensions-assume-record-high-investment-returns/#4c4547a46ec5.

Brown JR, Pollet JM and Weisbenner SJ (2015) The in-state equity bias of state pension plans. NBER Working Paper No. 21020. National Bureau of Economic Research: Cambridge, MA.

Chingos MW and West MR (2015) Which teachers choose a defined contribution pension plan? Evidence from the Florida retirement system. *Education Finance and Policy* **10**, 193–222.

Clark RL, Hanson E and Mitchell OS (2016) Lessons for public pensions from Utah's move to pension choice. *Journal of Pension Economics and Finance* **15**, 285–310.

Clemens J and Cutler DM (2014) Who pays for public employee health costs? *Journal of Health Economics* **38**, 65–76.

Costrell RM (2015) School Pension Costs Have Doubled Over the Last Decade, Now Top $1,000 per Pupil Nationally. Available at http://www.teacherpensions.org/blog/school-pension-costs-have-doubled-over-last-decade-now-top-1000-pupil-nationally.

[28]We are not aware of any comparable analysis of effect heterogeneity in the literature on health benefit costs. The data panels used by Clemens and Cutler (2014) and Lubotsky and Olson (2015) end in 2007 and 2008, respectively. The data panel used by Anand (2017) ends in 2010 but is predominantly pre-recession, and she does not explore the potential for differential crowd-out before and after the recession.

[29]As of 2013, the following states enrolled new entrants into a sub-plan/tier that was less generous than the plan in which some other more experienced teachers were enrolled: AL, AZ, DE, IL, IN, KS, KY, MS, NJ, NM, NY, ND, OR, SC, PA, UT, VT, WA, WI, and WY (Backes et al., 2016).

Costrell RM and Podgursky M (2010) Distribution of benefits in teacher retirement systems and their implications for mobility. *Education Finance and Policy* **5**, 519–557.

Even WE and Macpherson DA (2014) What do unions do to pension performance? *Economic Inquiry* **52**, 1173–1189.

Fitzpatrick M (2015) How much do public school teachers value their pension benefits? *American Economic Journal: Economic Policy* **7**, 165–188.

Friedberg L and Webb A (2005) Retirement and the evolution of pension structure. *Journal of Human Resources* **40**, 281–308.

Glaeser EL and Ponzetto G (2014) Shrouded costs of government: the political economy of state and local public pensions. *Journal of Public Economics* **116**, 89–105.

Goldhaber D and Grout C (2016) Which plan to choose? The determinants of pension system choice for public school teachers. *Journal of Pension Economics and Finance* **15**, 30–54.

Kan L and Aldeman C (2014) *Uncovered: Social Security, Retirement Uncertainty, and 1 Million Teachers*. Policy Report. Washington, DC: Bellwether Education Partners.

Koedel C and Podgursky M (2016) Teacher pensions. In Hanushek EA, Machin S and Woessmann L (eds), Handbook of the Economics of Education, Vol. **5**. Amsterdam: Elsevier Science, pp. 281–304.

Koedel C, Ni S and Podgursky M (2013) The school administrator payoff from teacher pensions. *Education Next* **13**, 8–13.

Koedel C, Ni S and Podgursky M (2014) Who benefits from pension enhancements? *Education Finance and Policy* **9**, 165–192.

Lubotsky D and Olson CA (2015) Premium copayments and the trade-off between wages and employer-provided health insurance. *Journal of Health Economics* **44**, 63–79.

McGee J and Winters MA (2013) Better Pay, Fairer Pensions: Reforming Teacher Compensation. *Center for State and Local Leadership. The Manhattan Institute*.

Munnell AH, Aubry J-P and Cafarelli M (2015) How Did State/Local Plans Become Underfunded? *State and Local Pension Plans* 42.

Novy-Marx R and Rauh J (2009) The liabilities and risks of state-sponsored pension plans. *Journal of Economic Perspectives* **23**, 191–210.

Novy-Marx R and Rauh J (2011) Public pension promises: how big are they and what are they worth? *Journal of Finance* **66**, 1207–1245.

Novy-Marx R and Rauh J (2014) The revenue demands of public employee pension promises. *American Economic Journal: Economic Policy* **6**, 193–229.

Petrella D and Pearson R (2019) Gov. J.B. Pritzker's administration considering selling state buildings and land to pay for pensions. *Chicago Tribune* (02.14.2019).

Public Plans Database (2001–2015) Center for Retirement Research at Boston College, Center for State and Local Government Excellence, and National Association of State Retirement Administrators. Available at http://publicplansdata.org/public-plans-database/download-full-data-set/.

Schmitz K (2016) Are Teacher Salaries Flat because of Changing Workforce Demographics? Teacher Pensions Blog (07.06.2016). Available at http://www.teacherpensions.org/blog/are-teacher-salaries-flat-because-changing-workforce-demographics.

Snyder TD, De Brey C and Dillow SA (2018) *Digest of Education Statistics 2016*, NCES 2017–094. Washington, DC: National Center for Education Statistics.

Tucker J (2017) California Schools May Face Cuts Amid Skyrocketing Pension Costs. *San Francisco Chronicle* (01.23.2017). Available at http://www.sfchronicle.com/bayarea/article/California-schools-may-face-cuts-amid-10873046.php.

Weingarten R (2017) Flawed Bellwether Education Report on Teacher Pensions. American Federation of Teachers Memorandum (05.22.2017). Available at https://www.aft.org/sites/default/files/memo_bellwether-report_052217.pdf.

Yin Y and Boyd D (2018) Analyzing the Interplay Between Public-Pension Finances and Governmental Finances: Lessons from Linking an Economic Model to a Pension Fund Model. Working paper prepared for the Brookings Municipal Finance Conference (07.17.2018).

Cite this article: Kim D, Koedel C, Xiang PB (2021). The trade-off between pension costs and salary expenditures in the public sector. *Journal of Pension Economics and Finance* **20**, 151–168. https://doi.org/10.1017/S1474747219000362