Pensions in the Trenches:
How Pension Costs are Affecting U.S. Local Government

Sarah F. Anzia
Associate Professor
Goldman School of Public Policy
Department of Political Science
University of California, Berkeley
sanzia@berkeley.edu

July 2019

Abstract: Are U.S. local governments experiencing dramatic increases in their public-employee pension costs? Are pension costs higher in places with strong public-sector unions and collective bargaining? And are rising pension costs crowding out local government services, as some experts claim? These questions are of fundamental importance to both scholars of American government and policy practitioners, but no existing studies—nor the datasets they rely on—provide answers to them. This paper analyzes a new dataset of the annual pension expenditures of over 400 municipalities and counties across the U.S. from 2005 to 2016. I find that pension expenditures are rising almost everywhere, and that the increases are larger in states with stronger unions. On average, local governments are not responding to rising pension costs by increasing revenue; they are instead shrinking their workforces and reducing capital outlays—patterns that again are more pronounced in places with strong unions and collective bargaining.

This project was generously funded by the Laura and John Arnold Foundation, the Berkeley Institute for the Future of Young Americans, and the Institute for Research on Labor and Employment. The author thanks Natalie Ahn, Brendan Rowan, Carson Whitelemons, Kimberly Rubens, Anna Duning, Lauren Finke, Brad Kent, Tim Tsai, Andrea Lynn, Devan Shea, and Annie Benn for excellent research assistance.

Over the last few decades, state and local policymakers have enacted changes to make government employees' pension benefits more generous, and they have also consistently underfunded those pensions—setting aside far too little money to pay for them. As a result, public employee pensions in the United States are underfunded by somewhere between $1.28 trillion and $3.4 trillion, depending on the estimates used (Pew Charitable Trusts, 2019; Rauh, 2016). Because pension benefits are legally guaranteed, someone has to make up for these shortfalls, and many experts claim that state and local governments are now feeling the painful consequences in the form of rapidly rising pension expenditures. As a 2011 *New York Times* article put it, because of rising pension costs, "many Americans may be forced to rethink what government means at the state and local level" (Walsh, 2011).

Based on a few state and local governments that have received a great deal of media attention, it is clear that pension costs *have* risen dramatically in some places—and that those places have struggled to cope with the fiscal pressure. In cities like San Jose, San Diego, Chicago, Detroit, and Stockton, skyrocketing pension costs have forced government officials to make painful decisions about future pension benefit levels, government service provision, and taxes. Some say that the events unfolding in these cities are part of a much larger and permanent trend—one affecting state and local governments almost everywhere—and that government employees and their unions have played a key contributing role (Anzia and Moe, 2017, 2019; DiSalvo, 2015; DiSalvo and Kucik, 2017; Kiewiet and McCubbins, 2014). Others argue that claims of widespread pension-induced fiscal crisis are exaggerated and that there is no relationship between pension costs and strong unions (Munnell, 2012; Munnell et al., 2013).

Which of these conclusions is closer to the mark? Given the state of the literature, it is impossible to know. There is existing research on public employee pensions, but none of those

studies—nor the datasets they rely on—allow evaluation of how much local governments are spending on pensions, the degree to which those expenditures are rising, or how local governments are responding to any pension cost increases. There has also been little assessment of how and why local governments' pension costs vary from place to place. The literature on the effects of public-sector unions focuses mainly on salaries and wages (e.g., Anzia and Moe, 2015; Hoxby, 1996; Lovenheim, 2009; Paglayan, 2019), and the few studies that do analyze retirement benefits do not test whether local pension costs are higher or growing faster in places with stronger unions (Anzia and Moe, 2017, 2019; DiSalvo and Kucik, 2017; Kiewiet, 2010; Munnell, 2012). As it stands, then, we have little sense of how much local governments spend on pensions, whether those costs are higher in places with public-sector unions, how local pension costs have changed over time, or what governments are doing to respond to any such changes.

These questions are basic, but they are of fundamental importance for scholars of American government because they are ultimately questions about what local government is, what it does, and how that is changing as a result of legacy costs. The nation's nearly 90,000 local governments play a critical role in government service provision: they spend roughly a quarter of all public money in the United States, provide critical services such as public education and public safety, and are responsible for local infrastructure like sewers and roads. A long tradition of local political economy research attempts to explain variation in local service provision, but it has so far emphasized the importance of local political institutions, partisanship, ideology, and demographics (e.g., Trounstine, 2018; Alesina et al., 1999; de Benedictis-Kessner and Warshaw, 2016; Tausanovitch and Warshaw, 2014; Gerber and Hopkins, 2011)—and has yet to examine the potentially important role of local retirement costs and public-sector unions.

Questions about local pension costs are also important for policymakers, policy practitioners, government employees, and citizens more broadly, because up to this point, the national debate about public pensions has focused on seemingly abstract quantities like unfunded liabilities and actuarial assumptions. While these quantities are central components of the overall problem, they are also technical, debatable, and hard to understand—making it difficult for most to engage meaningfully on the issue. By contrast, the questions I pose here are simple. Rather than engaging in the debate about what the discount rate should be, or how big unfunded liabilities are, my focus is on what is happening in local governments and how that is affecting the citizens they serve—advancing the national conversation about pensions in a way that is easier to appreciate and understand.

This paper begins to answer these questions through a multi-step analysis of a new dataset of the annual pension expenditures of over 400 municipal and county governments from 2005 to 2016, which I hand-collected from the cities' and counties' annual financial statements. This dataset is unlike any that existed before because it tracks actual local government pension expenditures over time, and not just in the largest cities or in the cities with the biggest problems, but instead in a large, diverse set of cities and counties across the country. With these new data, we can see for the first time what American cities and counties are spending on their employees' pensions, how that has changed over a twelve-year period, and whether pension expenditure levels and growth differ systematically in places with stronger and weaker public-sector unions. In addition, by connecting these local pension expenditure data with U.S. Census data on local government employment and finances, I evaluate whether growing pension expenditures are associated with increased local revenue or cuts to public service provision, shedding light on how pensions are affecting how those local governments operate.

Three main sets of findings emerge from the analysis. First, cities and counties vary dramatically in how much they spend on retirement benefits, and on a per-employee basis, local governments in states with strong public-sector unions and mandatory collective bargaining do spend significantly more than those in weaker-union, non-bargaining states. Second, between 2005 and 2016, city and county pension costs rose in real terms almost everywhere, both per employee and as a share of general revenue. Again, though, there is substantial variation in the extent of that pension cost growth, and strong unions and collective bargaining are associated with larger increases. Third, local governments on average are not responding to rising pension costs by increasing revenue; they are instead shrinking their workforces and reducing capital outlays—especially governments in states with strong unions and collective bargaining. Stepping back, then, the picture that emerges is one of rising local pension costs and the crowding out of government service provision—but with substantial variation in the extent of that transformation.

Background and Literature

Approximately 14 million people work full-time for U.S. state and local government, and almost all of them are eligible for a defined benefit retirement plan (a traditional pension). Thus, unlike in the private sector, where defined contribution plans are most common, state and local government employees receive a defined benefit in retirement for as long as they live, equal to some fraction of their final average salary times the number of years they worked for the government. Most state and local employees are enrolled in large, state-operated pension plans such as CalPERS in California and OPERS in Ohio, but many local governments operate their own plans. In principle, the model for funding pensions is straightforward: they are supposed to be prefunded, with government employers and employees setting aside funds to pay for the retirement benefits earned each year.

Yet today, most state and local pension funds do not have sufficient assets to cover the retirement benefits that have been promised. Collectively, the nation's public-employee pensions are underfunded by somewhere between $1.28 and $3.4 trillion—estimates that depend on the assumptions used to value liabilities and assets (e.g., Pew Charitable Trusts, 2019; Rauh, 2016).

A number of decisions at both the state and local levels have contributed to this massive shortfall. First, over the years state and local officials have enacted changes to make pension benefits more generous and thus more expensive (DiSalvo, 2015; Anzia and Moe, 2017), such as increasing the benefit formula's multiplier, reducing the retirement age, allowing employees to purchase service credit for years not actually worked, and building in automatic cost-of-living adjustments. These changes have had significant and long-lasting effects, because in most states, pension benefits can only be reduced for future government hires—not for current employees, even for years those employees have not yet worked.

Governments have also consistently underfunded their pensions, setting aside far too little money to pay for the benefits they have promised. Many different kinds of decisions and events have contributed to the underfunding, including strategically-chosen actuarial assumptions that make pension liabilities look smaller than they actually are (Novy-Marx and Rauh, 2009, 2011), failure to pay the annual amounts required for full funding (Anzia and Moe, 2019), politically-motivated investment decisions (Andonov et al., 2018), and the decline in asset values brought by the Great Recession. Regardless of how poorly-funded pensions are, however, the benefits are legally guaranteed, and therefore someone has to make up for the shortfalls.

There is strong reason to expect that these trends are affecting local government budgets, but the existing research literature has done little to study what local governments are experiencing or how they have responded. It has instead focused mostly on outcomes related to

the largest state and local pension *plans*, such as CalPERS and OPERS (e.g., Thom, 2013; Stalebrink, 2014; Mitchell and Smith, 1994). One line of work, for example, attempts to explain variation in plans' funding ratios and investment performance (e.g., Andonov et al., 2018). Another focuses on estimating what public pension liabilities are actually worth given different sets of actuarial assumptions (Novy-Marx and Rauh, 2009, 2011, 2014). Plan-level outcomes presumably do have effects on the local governments that participate in those plans, but so far the research literature has not directly studied those effects at the local government level.

Scholars also debate how public-sector unions and collective bargaining influence pension benefits and funding levels, but here, too, the empirical literature is thin. Most research on the fiscal effects of public-sector unions examines salaries and wages: it finds that local governments with collective bargaining and strong public-sector unions do pay higher salaries, but there is debate about whether that relationship is causal (Anzia and Moe, 2015; Hoxby, 1996; Lovenheim, 1999; Paglayan, 2019). DiSalvo and Kucik (2017) evaluate states' other post-employment benefit (OPEB) liabilities and find that stronger unions are associated with higher OPEB costs, but they do not analyze pension costs or any retirement obligations at the local level. On pensions specifically, Anzia and Moe (2017, 2019) provide evidence that strong public-sector unions make state legislators less favorable toward benefit reductions and are associated with less responsible funding decisions, but they do not look at how pension costs vary by union strength, nor do they examine budgetary effects in local governments. Meanwhile, at the level of the pension plan, Munnell (2012) finds that the "normal cost" of pensions as a share of payroll—where "normal cost" is the value of pension benefits earned by employees in a year—has no significant relationship to state collective bargaining laws.

Why haven't scholars studied how pension costs are affecting local government or whether those costs are different in places with strong public-sector unions? The answer, almost certainly, is that there aren't any readily available data on pension costs at the level of the local government. Nearly all of the aforementioned empirical work relies on the Public Plans Database developed by Boston College's Center for Retirement Research, which documents each state and large local plan's funded ratios, actuarial assumptions, required contributions, and more. Yet these plan-level data do not tell us about the pension expenditures of particular governments, most of which contribute to multiple pension plans—typically at least one state-operated plan and often one or more locally-administered plans. The problem is therefore a mismatch between the unit of analysis in available datasets—the pension *plan*—and the unit of analysis needed to study what local governments are experiencing—which is the local *government*.[1]

Because of the lack of data, we do not actually know how rapid or pervasive local pension cost increases have been, whether those costs vary with the strength of public-sector unions and collective bargaining, how governments are responding to any such pension cost changes—or, for that matter, how much local governments even spend on pensions. Some scholars argue that pension costs are a heavy burden for local governments, that increases in pension costs are significant and widespread, and that local governments throughout the country are facing difficult budgetary choices as a result, but they rely on state- or plan-level data or examples from a few cities to draw that conclusion (Kiewiet and McCubbins, 2014; DiSalvo,

[1] Using pension plan-level data to estimate how much local governments spend or should be spending on pensions is a difficult and imperfect exercise. See Munnell and Aubry (2016) for an example of this approach.

2015). Others argue that the experiences of cities like New York and San Jose might not be typical of other local governments and that most local governments are not experiencing such fiscal stress (e.g., Munnell et al., 2013; Munnell and Aubry, 2016). Without data on what local governments spend on pensions, we cannot know which is closer to the mark.

Data

To evaluate how local governments' pension costs have changed over time, I assembled a new dataset. I started by selecting 236 municipal governments and 239 county governments from among those that appear in the U.S. Census Bureau's Survey of Governments Finance and Employment files for most years between 2005 and 2016.[2] I then set out to collect the comprehensive annual financial reports (CAFRs) of each of those 475 local governments for that same twelve-year period. Most local governments had at least some CAFRs on their websites, typically for the most recent years. For city- and county-years for which the CAFRs were not available online, I requested the documents from the local governments, filing public information requests where necessary. I was able to obtain at least some years' CAFRs for 442 local governments, including 227 municipalities and 215 counties.

CAFR formats vary, but nearly all of them provide information about the retirement plans to which the government contributed on behalf of its employees (most cities and counties have more than one plan), each plan's Annual Pension Cost or Actuarially Required Contribution, and the amount of the employer contribution. Therefore, the most important piece of information I drew from the CAFRs was the amount the city or county contributed to each of its employee retirement plans in that year.[3] I included contributions to defined contribution (DC) plans as well

[2] See the online appendix for details.

[3] I provide a detailed account of the data collection and coding in the online appendix.

as defined benefit (DB) plans, although DC plans in these local governments are rare and typically make up a small share of total retirement contributions. A small number of governments also fund their OPEB from their pension fund contributions. I subtracted out funds going to OPEB whenever possible, but I was not able to do this for a small number of plans, and in those cases, the pension contribution amounts include some OPEB expenditures.

Three other features of the data collection are worth highlighting. First, while my goal was to collect cities' and counties' total pension contributions, including any amount of the employees' contributions paid by the local government (Employer-Paid Member Contributions, or EPMC), most CAFRs did not clearly and consistently report whether the local government was paying EPMC or, if it was, how much. Therefore, the retirement contributions discussed below do not include EPMC. They also do not include contributions the local governments made using revenue from pension obligation bonds (POBs) or any interest paid on those bonds, even though both can be substantial.[4] For some cities and counties, then, the contribution amounts discussed below are smaller than they would be if POB interest payments and EPMC could have been included. Third, the dataset tracks what governments actually paid toward retirement benefits—not what they should be paying. Most work on public pensions focuses on underfunding and what governments should be contributing, but my focus is on whether pension expenditures have risen over time and how that is affecting local government. The more

[4] When local governments issue POBs, they typically use the revenue from the bonds to make a large, non-recurring pension contribution. I subtracted out those POB funds from the contribution amounts.

appropriate measure for this study is what local governments are *actually* spending on pensions, because those are the numbers affecting their budgets.

This dataset is unlike any that existed before, and it is uniquely suited to the task of assessing the on-the-ground experiences of American cities and counties. For the analysis to follow, I summed the retirement expenditures for all plans in each city- and county-year. In total, the dataset has 5,085 annual pension expenditure observations from 442 unique governments, spanning all 50 states plus Washington, DC, and including localities as small as Parsons, West Virginia, and Eureka County, Nevada, each with fewer than 2,000 residents, and as large as New York City and Los Angeles County. For 375 local governments, the dataset includes pension expenditure information for all twelve years from 2005 to 2016, and for the remaining 67, it includes pension expenditures for some. For the first time, then, these data allow us to see what American local governments are spending on their employees' retirement benefits and how that has changed in recent years.

Local Pension Expenditures, 2005-2016

As a first and simple step, I analyze how pension expenditures vary across the 442 cities and counties in the dataset. After adjusting each year's total pension expenditures for inflation (to 2016 dollars), I calculate two variables for each local government and year: total pension expenditures as a share of general revenue, and total pension expenditures per full-time equivalent employee.[5] While both variables are of interest, the second is a clearer measure of pension-related fiscal pressure. To see why, consider that a local government's pension costs are a direct function of how many employees it has: if a city hires more employees, for example, its

[5] Local government general revenue comes from the U.S. Census Surveys on Finance, and full-time equivalent employment comes from the U.S. Census Surveys on Employment.

total pension contributions should automatically increase because it is contributing on behalf of more people. Thus, pension expenditures as a share of general revenue could be higher in some cities and counties just because they have more employees, and they could be increasing within a city or county just because the government is expanding its workforce—which probably is not a sign of fiscal stress. Pension costs per employee, by contrast, takes into account the size of the workforce. It captures the amount the local government is setting aside for pensions per worker—and should generally be higher in governments and years where pension benefits are more generous or where governments are making up for larger funding shortfalls.

Table 1 presents descriptive statistics for these two variables as of the year 2007—an early year in the dataset and one for which I have pension expenditure data and U.S. Census employment and finance data for most of the 442 local governments. That year, the median local government in the dataset spent $4,901 per full-time equivalent employee on pensions, amounting to 3.1% of its general revenue, but there was tremendous variation across cities and counties. In the top 25%, pension expenditures were more than $7,726 per employee and 4.7% of general revenue. And there is substantial right skew in the distributions of these variables. The top 5% of cities and counties were spending more than $15,183 per employee and more than 7.7% of general revenue on pensions. In the top 1% of these local governments, pension expenditures in 2007 were more than 10% of general revenue and over $21,728 per employee. Clearly, even as of 2007, U.S. cities and counties varied considerably in their pension costs.

Are local pension costs higher in the presence of strong public-sector unions and collective bargaining? While this might seem a controversial question—one on which scholars have made divergent arguments (e.g., Munnell, 2012; DiSalvo, 2015)—there is one straightforward reason to expect the answer to be yes: public employees' salaries tend to be

higher in cities with unionized employees and collective bargaining (Anzia and Moe, 2015), and employees' pension benefits are a function of their final average salaries.[6] Public-sector unions might also be more successful in pushing for a higher benefits in other ways, such as by increasing the multiplier. However, if pension *underfunding* is also more pronounced in states with stronger unions (Anzia and Moe, 2019), then that should affect the relationship as well—and the direction of any such effect is theoretically ambiguous. On the one hand, contributions themselves would tend to reflect a greater tendency to underfund in heavily unionized states (suggesting they should perhaps be lower), but on the other hand, cities and counties in heavily unionized states would eventually be called on to contribute more to make up for greater accumulated shortfalls (suggesting they should perhaps be higher). It is also not clear whether to expect local governments with strong unions to spend a larger share of their budgets on pensions: they may dedicate larger sums to employee-related costs, but they may also have larger budgets overall. It is an empirical question, then, whether local governments with stronger unions have higher pension expenditures.

To test this, I use OLS to regress local pension costs as of 2007 on the proportion of state and local government employees who are members of unions in each local government's home state.[7] (There are no modern data on public-sector union membership at the level of the local

[6] Examining the normal cost of pensions as a proportion of payroll, as Munnell (2012) does, would mask such a difference if both pension costs and payroll are higher in places with strong unions.

[7] These figures are Anzia and Moe's calculations using Current Population Survey data from 2000 to 2010; the variable is constant within states for all years of the dataset.

government.) In addition, I include log per capita income in the local jurisdiction, expecting pension costs to be higher in places with higher cost of living. Because union membership tends to be higher in more liberal places, and more liberal cities and counties might be more supportive of greater pension spending, I also include local presidential vote share for Barack Obama in 2008.[8] DC plans are often touted as less expensive for governments than DB plans, so I include a variable equal to the proportion of the local government's annual retirement expenditures that went to DC plans. Munnell (2012) proposes that local governments with larger numbers of plans or more locally-operated plans may have lower funding levels, so I control for the number of plans in the locality as well as the share of pension expenditures that went to locally-operated plans. Finally, I include the proportion of the annual pension expenditures that also went to paying for OPEB, expecting it might be associated with higher overall contribution amounts.

While these are seven variables for which there is some reason to expect a relationship to local pension costs, in all models I also add a series of variables that have been emphasized in the literature on local politics: the log of the city or county's population and the proportions of local residents who are living in urban areas, who are homeowners, and who are black, Asian, and Hispanic. To explore whether there are average pension cost differences between cities and counties, I include a dummy variable for counties. I cluster the standard errors by state.

Column 1 of Table 2 presents the estimates for the model of local pension expenditures as a share of general revenue in 2007. The coefficient on union membership is statistically indistinguishable from zero, showing that cities and counties in states with stronger unions did

[8] City-level presidential election returns were not available for a few cities, so for them, this variable equals presidential vote in the parent county.

not spend a significantly larger share of general revenue on pensions than local governments in states with weaker unions. I do find, however, that greater use of DC plans is associated with spending a smaller share of general revenue on pensions, and that more urban areas and cities spend a greater share on pensions than less urban areas and counties. None of the other coefficients, however, are statistically significant.

In column 2, I model logged pension costs per employee in 2007 (in thousands of 2016 dollars),[9] and there I estimate a clear, positive association with the strength of public-sector unions. On average, a 10-percentage-point increase in public-sector union membership is associated with 5.6% increase in pension spending per employee as of 2007. To put this in context, a 5.6% increase from the median of $4,901 (see Table 1) amounts to $276 per government employee. Comparing a state like California with 62% union membership to Colorado with 26% union membership, the model predicts a 22% increase in pension spending per employee, which would be an increase of $1,069 from the median. A few of the other variables are also significantly associated with higher pension costs per employee, including per capita income, larger population, more urban population, larger shares of Asian and Hispanic residents,[10] and being a municipal as opposed to a county government. Interestingly, more liberal

[9] I log per-employee pension costs to reduce the substantial right-skew of this variable. In taking the log, I add one to all values to preserve the few cases where contributions equaled zero.

[10] This is a California effect: when I add an indicator for whether the local government is in California, these two coefficients are no longer statistically significant.

cities do not pay higher pension costs per employee,[11] nor do local governments with more plans or more locally-operated plans.

Because public-sector union membership is so highly correlated with state laws requiring collective bargaining (see Saltzman, 1985, 1988; Moe, 2011),[12] there is little to be gained by estimating separate effects of union strength and collective bargaining in this model.[13] In column 3, however, I replace the union membership variable with an indicator for whether the local government is in a state that requires government employers to collectively bargain with their employees if those employees form a union. Using the public-sector labor law dataset compiled by Anzia and Moe (2016), I code each city and county as 1 if police, firefighters, and other local employees in the state are covered by a duty-to-bargain law and 0 if any of them are not. The estimated coefficient on *Collective bargaining* is 0.194, suggesting that as of 2007, shifting from a non-duty-to-bargain state to a duty-to-bargain state is associated with a 21% increase in pension expenditures per employee. Thus, cities and counties with strong unions and collective bargaining do spend significantly more per employee on pensions.

Next I evaluate how pension costs have changed over time within this set of local governments. As a starting point, I calculate percent growth in total inflation-adjusted pension expenditures from 2005 to 2016 for the 359 cities and counties for which I have data for both

[11] The same is true when I use the Tausanovitch and Warshaw (2014) ideology scores instead of presidential vote; see the online appendix.

[12] The correlation between the two variables in this dataset is 0.8.

[13] When I include both the collective bargaining indicator and the unionization measure in the model, both are positive but statistically insignificant.

years.[14] The distribution of that variable is shown in the top left panel of Figure 1. Strikingly, the number is positive for almost all cities and counties: 88% experienced increases in their pension expenditures from 2005 to 2016. The median within-government change over twelve years is 56%—quite large—but most notable is the long right tail of the distribution. In 26% of the cities and counties, pension costs more than doubled in twelve years. In the top 10%, they grew by more than 189%. And in a non-negligible number of cities and counties, the pension cost changes from 2005 to 2016 were very large, such as in Lacey, Washington (388%) and Merced County, California (539%).[15]

This first figure only tells us so much, however, because rising pension costs could simply be a sign of a growing city budget or public-sector workforce. In the top-right panel, therefore, I present the distribution of the change in pension costs as a share of general revenue from 2005 to 2016.[16] It shows that pensions have grown as a share of revenue in 75% of the cities and counties, with a median change of 0.008 (or 0.8 percentage points). More notable, again, is the right tail: the top 25% of the cities and counties saw pension expenditures consume an additional 2.1% of general revenue or more, and the top 10% had pensions absorb an additional 3.7%.

---

[14] This number excludes a small number of cities and counties for which I have pension expenditure data for both 2005 and 2016 but that, due to changes in reporting or missing information on one or more plans, the 2005 and 2016 figures are not comparable.

[15] This plot excludes Monmouth County, New Jersey, whose pension expenditures increased by 2815% from 2005 to 2016.

[16] This, too, excludes an outlier: Plymouth County, Massachusetts.

In the bottom two panels, I show the within-government change in pension expenditures per local government employee, first for 2005 to 2016 (which only includes 269 cities and counties, mostly due to missing employment data for 2005), and then on the right for 2007 to 2016 (which includes 366 cities and counties).[17] Both show that the vast majority of cities and counties have seen increases in pension expenditures per employee. Focusing on the bottom-left, the median within-government change from 2005 to 2016 was $1,419 per employee, and in 25% of the cities and counties, per-employee pension expenditures increased by $3,542 or more. In a now-familiar pattern, the increases in per-employee pension expenditures have been very large in some places. In the top 10% of cases, pension costs have increased by more than $7,124 per employee, and in the top 5%, they have increased by more than $11,321 per employee.

Regardless of how it is measured, then, most cities and counties have experienced growth in their pension costs, but there is considerable variation in the extent of that growth. In Table 3, I explore characteristics of local governments and their retirement plans that might predict more or less growth in pension costs, regressing the last three measures shown in Figure 1 (a single observation for each government) on the same set of independent variables as in Table 2.

Column 1 presents the estimates for the model of change in pension expenditures as a share of general revenue. It shows that the increases have tended to be bigger in cities and counties with higher public-sector union membership: an increase of 10 percentage points in union membership is associated with a 0.19-percentage-point greater increase in the proportion of general revenue going to pension expenditures between 2005 and 2016. In column 2, I model

[17] For presentation purposes, these figures exclude a few outliers: Pembroke Pines, Florida, in both, and also San Jose, California, and Joliet, Illinois, on the bottom right.

the change in pension costs per full-time equivalent employee from 2005 to 2016 (in thousands of 2016 dollars), and I again find a statistically significant positive relationship between public-sector union strength and growth in pension costs. Here, a 10-percentage-point increase in union membership is associated with additional pension cost growth of $395 per employee over this twelve-year period. I find the same pattern in column 3, where I replace the union membership variable with the indicator for collective bargaining, and in column 4, where I model growth in pension costs per employee from 2007 to 2016. On average, cities and counties in duty-to-bargain states saw their pension costs grow by an additional $884 per employee from 2005 to 2016 compared to those in non-bargaining states. And focusing on 2007-to-2016 growth, a 10-point increase in union membership is associated with greater pension cost growth of $300 per employee.

Two other clear findings that emerge are that pension costs have grown more slowly in localities that rely more on DC plans and more rapidly in places with more locally-operated plans. In column 4, for example, I find that cities and counties with 50% spending on DC plans had pension cost growth that was $1,424-per-employee less from 2007 to 2016 than those with 0% spending on DC plans. The same model estimates suggest that local governments with 50% spending on locally-administered plans had an additional $1,234 increase in pension expenditures per employee than governments with all state-operated plans.

In sum, pension costs have indeed increased in most cities and counties since the mid-2000s—in absolute terms, as a share of general revenue, and per employee. But there is considerable variation in the extent of that growth. While the median local government saw pension costs grow by about $1,400 per employee over this twelve-year period, many cities and counties experienced increases much, much larger than that. The upshot is that cities like San

Jose and Chicago—with their staggering pension cost increases—are by no means representative of all local governments, but they are also not alone. Many other cities and counties are experiencing the same thing—and those that are do tend to be in states with strong public-sector unions and mandatory collective bargaining.

How Are Pension Costs Affecting Local Government?

Given that pension costs have increased in most local governments—in some places substantially so—an important next step is to analyze how local governments are responding to those changes. To do this, I turn to the full panel dataset of city and county pension expenditures from 2005 to 2016 and evaluate whether larger pension cost increases are associated with trends in local finances and employment. I focus in particular on whether rising pension costs are leading to increases in revenue, decreases in employment, or decreases in spending on items other than retirement benefits.[18] Local governments might reasonably adjust course in different ways, and so my approach is evaluating whether there are discernable trends in cities' and counties' responses—and any clear links between pension costs and government operations.

I first explore whether cities and counties tend to cope with rising pension costs by increasing revenue. I model two dependent variables, both of which are from the U.S. Census Survey of Governments Finance files for 2005 to 2016: the log of total general revenue per

[18] Another option local governments may have is to issue POBs. My coding of CAFRs identified 38 cities and counties that either issued or were paying off POBs (14 of which are in California). However, only 13 of the 38 issued POBs during this study period—too few to carry out analysis of whether greater pension cost increases are associated with use of POBs.

capita and the log of total own-source general revenue per capita, adjusted to 2016 dollars.[19] While general revenue (which includes intergovernmental revenue) better captures the total revenue cities and counties have at their disposal, own-source general revenue may more clearly reflect local government actions to increase revenue in response to rising pension costs.[20] Throughout, the main independent variable of interest is the log of pension expenditures per full-time equivalent employee.[21]

Most of the variation in per capita general revenue is across rather than within cities and counties, but here I am focused on how cities and counties might be *changing* their general revenue in response to rising pension costs. I therefore model the general revenue variables with OLS and fixed effects for each city and county, which partial out the influence of any time-constant characteristics of the local governments that lead them to have higher or lower general revenue and pension expenditures. I also include year fixed effects because there are likely secular trends that affect pension costs and general revenue in all cities. During the Great Recession, for example, required pension expenditures increased because of the overall decline in fund asset values, and at the same time, city and county revenues dropped. More generally, though, the analysis above made it clear that some cities and counties have had much greater

[19] I log-transform the data because the distributions of both general revenue per capita and own-source general revenue per capita are right skewed.

[20] Unfortunately, I know of no existing data on local government decisions about tax rates, assessments, or charges that covers all of the governments in this dataset.

[21] As I discussed earlier, pension expenditures per employee is the better measure of the extent to which cities and counties are under fiscal stress because of pensions.

increases in pension expenditures per employee than others, and including year fixed effects allows me to test whether those greater-than-average increases are associated with greater-than-average increases in general revenue.

I also lag the pension expenditure variable by one year so that I am estimating the effect of pension costs in year *t-1* on general revenue levels in year *t*. This models government decision-making in a realistic way; presumably officials make decisions about next year's budget based on what they observe of this year's. An insignificant coefficient on per-employee pension costs would therefore imply that on average, greater increases in pension expenditures per employee the year prior are *not* associated with greater increases in general revenue per capita. A positive coefficient would imply that greater increases in per-employee pension costs *are* associated with larger revenue increases—and would be evidence that the average local government is increasing revenue to make room for rising pension costs.

Finally, because there might be changes in the local jurisdiction that affect general revenue and may be correlated with pension cost increases, I include the time-varying local demographic variables from earlier: log per capita income, log population, percent urban, percent homeowners, and percent black, Asian, and Hispanic.[22]

[22] The other independent variables from Tables 2 and 3 are either constant within local governments or were included earlier for their potential to explain pension cost levels and growth—which is not the focus here. I lose a small number of observations for a few reasons: because pension costs for some city- and county-years are not comparable to other years within the same government, because of clear errors in the finance and employment data, or because of extreme changes in pension expenditures for a single year. See online appendix for details.

The estimates from this model are shown in column 1 (general revenue per capita) and column 2 (own-source general revenue per capita) of Table 4. In both, the coefficients on pension costs per employee are close to zero and statistically insignificant. Certain other city- and county-level variables are related to growth in general revenue per capita, such as per capita income and population, but there is no evidence of a link between rising pension costs and increasing general revenue: in cities and counties that experience greater-than-average increases in pension costs per employee, the next year does not bring greater-than-average increases in revenue.

In columns 3 and 4 of Table 4, I consider whether cities and counties with strong unions and collective bargaining are more likely to respond to rising pension costs with revenue increases. After all, officials in those local governments might place greater priority on maintaining employment levels and employee compensation in the face of fiscal pressure, and strong unions and labor-friendly laws tend to be found in more liberal places—where increasing revenue might be more palatable. I test this by interacting the pension cost variable with the indicator for collective bargaining. The results in columns 3 and 4 show that cities and counties in states with mandatory bargaining are no more likely than those in non-bargaining states to respond to larger pension cost increases with revenue increases. Even in collective bargaining states, there is no relationship between pension costs and revenue increases.[23]

[23] I use the collective bargaining indicator rather than union membership for ease of interpretation. The results are substantively the same with union membership. When I interact the pension cost variable with Democratic presidential vote, I find that the relationship between pension costs and general revenue does not depend on how liberal the city or county is.

Next I carry out a series of tests to evaluate whether rising pension costs are crowding out public service provision. There are many possible ways of thinking about public service provision—including dimensions such as frequency, geographic scope, equity, or quality—but as a practical matter, many of these dimensions cannot easily be measured for a large number of cities and counties over multiple years. One feature of local government that *can* be measured is its employment levels. Employment is also a natural place to begin looking for pension-related service reductions. More so than at the national level, local government service provision is heavily dependent on the employees providing the services. And because a local government's pension costs are in part a function of its employment levels, officials confronting rising pension expenditures might find that reducing employment is an attractive option.

Modeling the relationship between pension costs and local employment is less straightforward than it might seem, however, because the independent variable of interest—pension expenditures per employee—itself has employment in the denominator. Lagging the pension cost variable by one year (as I did for the models of general revenue) helps to address the mechanical endogeneity of pension costs and employment in the same year—and again, it is a plausible model of government decision-making.

The dependent variables come from the U.S. Census Survey of Governments employment files, which have information on full- and part-time government employment and payroll for 92% of the city- and county-years in my pension costs dataset. Large cities naturally have more employees than small cities, so I divide each of the employment dependent variables (detailed below) by local government population in thousands. The distributions of these employment variables are skewed right, so I take the log of each one. I model them using the same approach as in Table 4, with local government local government fixed effects, year fixed

effects, and time-varying local demographics. The main focus is on log pension costs per employee: An insignificant coefficient would imply that on average, larger increases in per-employee pension expenditures are *not* associated with greater employment reductions the following year. A negative coefficient would imply that greater increases in per-employee pension costs *are* associated with larger employment reductions the next year—and would be evidence that pension costs are leading to service reductions.

In column 1 of Table 5, I present the results from a model of logged full-time equivalent employment per capita. The coefficient estimates suggest that rising pension expenditures have indeed led to an average reduction in public-sector employment: a 10% increase in pension expenditures per employee is associated with a 0.66% decrease in employment the following year. To get a sense of the magnitude of this effect, consider that the median increase in pension costs per employee from 2007 to 2016 was $1,203, and that is approximately a 25% increase from the 2007 median pension expenditure per employee ($4,901, see Table 1). The coefficient estimate in column 1 of Table 5 suggests that a 25% increase in pension costs is associated with a 1.63% decrease in local employment per capita. Given that the median local government in this dataset had 10.13 full-time equivalent employees per thousand residents as of 2007, a 1.63% decrease represents the loss of 0.17 employees per thousand residents—or for a city or county of 100,000 people, 17 of its employees. Naturally, of course, the model predicts larger employment losses for the cities and counties that experienced larger growth in pension costs.

Given that local governments are reducing employment in response to rising pension costs, there is good reason to expect that the cuts would be greater among full-time employees than part-time local employees, because part-time employees often are not eligible for pensions. I explore this in columns 2 and 3 of Table 5. In column 2, the dependent variable is the log of the

city's full-time employees per thousand residents. The coefficient on log per-employee pension expenditures is negative and statistically significant, suggesting that a 10% increase in pension costs is associated with a 0.68% reduction in full-time employment. When I instead model the log of the city's part-time employees per capita, in column 3, the coefficient on pension costs is statistically insignificant.[24] This suggests that growing pension costs are associated with declining numbers of full-time employees—not part-time employees.[25]

Do cities and counties in states with strong unions and collective bargaining respond similarly to those in weaker-union, non-bargaining states when confronted with a pension cost increase of similar magnitude? One possibility is that strong unions with collective bargaining can better stave off employment reductions than weaker unions. Another possibility is that officials in cities with strong unions and collective bargaining have fewer levers to keep costs down, such as limiting salary increases or reducing other fringe benefits like OPEB. To explore this, I model logged full-time employment per thousand residents, this time interacting the pension costs variable with the union and collective bargaining variables.

Column 4 of Table 5 presents the results of the model with the interaction between pension costs per employee and the indicator for collective bargaining. The coefficient on the

[24] The number of observations is lower due to a few local governments without part-time employees.

[25] Local governments could also reduce the hours worked by part-time employees even if the number of part-time employees stays the same. However, when I model logged part-time employee payroll in each government, which reflects total hours, I find no significant relationship with pension costs.

pension expenditures variable remains negative, but it is smaller than before and statistically insignificant. Thus, for the average local government in a state *without* mandatory bargaining for local employees, larger pension expenditure increases are *not* significantly associated with larger decreases in full-time employment. For cities and counties in duty-to-bargain law states, however, the effect is large and negative: combining the coefficients on log per-employee pension expenditures and its interaction with collective bargaining shows that a 10% increase in per-employee pension costs is associated with about a 0.79% decrease in full-time local government employment. In column 5, I run the same model with the union membership variable instead of collective bargaining, and the results are substantively the same. Thus, the relationship between rising pension costs and employment reductions is more pronounced in places with strong public-sector unions and collective bargaining.

As a final step in my analysis of local government employment, I explore whether certain kinds of employees are more affected by rising pension costs than others. I analyze cities and counties separately because they have different mixes of employees: municipalities rely heavily on police officers, firefighters, parks and recreation employees, and sanitation workers, while counties have larger concentrations in corrections, health, the judicial system, and highways. Even within these two types of government, however, there is considerable variation in the employee mix from place to place, especially in counties. One way to simplify the analysis, then, is to compare public safety employment with non-public safety employment. Both cities and counties have substantial public-safety functions: in the median city, 40% of all full-time city employees are in either police protection or fire protection (25% police and 14% fire), and in the median county, both corrections and police protection employees each make up over 9% of the county workforce. Also, police and fire protection employees (and sometimes correctional

officers) often have their own pension plans, sometimes with more generous benefits than non-safety workers, which might make them more vulnerable to pension-related employment cutbacks than non-safety workers. That said, public safety employees tend to be well-organized and active in city politics (Anzia and Moe, 2015), and it is possible that their political strength helps to insulate them from employment reductions.

I start with models of the full-time employment levels of three categories of city employees: police protection employees, fire protection employees, and non-safety, non-education employees. (Very few of these cities handle education, but in the cities that do, the education employees are a large share of the total city workforce, so I exclude them.) I continue with the model set-up from column 4 of Table 5, allowing the relationship between pension costs and employment to vary depending on whether the city is in a duty-to-bargain law state. The results are shown in Table 6. In column 1, I find that larger increases in pension expenditures per employee are indeed associated with greater reductions in police protection employees, but only in the states with mandatory bargaining: a 10% increase in per-employee pension expenditures is associated with a 0.49% reduction in full-time police protection employment per capita. For fire protection employees, modeled in column 2, I find a larger association, but again only in states with collective bargaining: a 10% increase in pension costs is associated with a decrease in fire protection employment of about 0.71%. For non-public-safety employees, including parks and recreation and sanitation, the marginal effect for collective bargaining cities is larger still; column 3 shows an average decrease of 0.96% in employment per capita for every 10% increase in pension costs. Thus, in municipal governments in states with duty-to-bargain laws, rising pension costs have a clear association with employment reductions in police protection, fire protection, and non-public-safety areas.

The results for counties show a different pattern. In column 4, I analyze police protection employment per thousand county residents, and I find that greater increases in pension costs are *not* associated with greater reductions in police protection employment—regardless of whether the county is in a state with mandatory bargaining. For county corrections employees, moreover, larger increases in pension costs *are* associated with greater employment reductions—but only for counties in states without collective bargaining. Among non-safety (and non-education) employees, however, the effects are similar to those in cities: a 10% increase in pension costs is associated with a 0.63% reduction in employment, but only for counties in states with collective bargaining. Thus, the general trend is one of employment reductions connected to rising pension costs, particularly in local governments that have mandatory collective bargaining, although the specifics of how those reductions are distributed are different for cities and counties.

In a final round of analysis, I look beyond crowd-out in the form of employment reductions and test for a link between rising pension costs and capital outlays.[26] This is an important dimension of local government activity because it relates to its investments in

[26] For many of the city and county spending variables in the U.S. Census Survey of Governments Finance files, it is not clear whether they include pension expenditures, and it would be problematic to analyze the effects of rising pension costs on expenditure variables that might include those very pension costs. One nice feature of both the employment figures and the capital outlays figures is that I can be confident they do not include pension expenditures.

construction and the purchase of land, equipment, and existing buildings.[27] In Table 7, I regress logged capital outlays per capita on the same variables as in the earlier tables, once again separating municipal and county governments.

The results for cities are shown in column 1, and column 2 shows the estimates for counties. For counties, I find that larger pension costs are not associated with greater reductions in capital outlays, but the estimates for municipal governments show a different pattern. For cities in states without collective bargaining laws, per-employee pension costs are associated with lower spending on capital projects, but the effect is not significant. For cities in states with bargaining laws, the effect is significant: a 10% increase in per-employee pension costs is associated with about a 3% decrease in capital outlays per capita. Thus, it appears that the increase in pension expenditures is not only linked to a decrease in city employment but also to a decline in spending on construction and equipment.

Conclusion

Nearly all scholarly work on public-employee pensions up to this point has focused on explaining variation in plan-level quantities such as funding ratios, unfunded liabilities, investment returns, and changes to benefit formulas. This body of research has shed considerable light on the political, economic, and actuarial decisions that have contributed to the underfunding problems that plague so many U.S. pension plans today. But on the issue of how local governments across the U.S. have been affected, scholars have advanced different views. Some

[27] The functional areas with the largest capital outlays in a typical city are sanitation and highways (including streets, toll highways, bridges, tunnels, ferries, street lighting, and snow removal).

claim that generous pension benefits combined with severe underfunding—both a result of public-sector union pressure, it is argued—are putting strain on local government budgets, and that widespread pension-induced fiscal pressure is forcing local governments to cut back on public services. Others maintain that the big problems are limited to a few governments that have been fiscally irresponsible, and that public-sector unions and collective bargaining have played little to no role. Until now, researchers haven't had the data necessary to assess which is closer to the mark.

This paper introduces a new dataset that allows us to evaluate for the first time where retirement expenditures tend to be higher, the extent to which they have increased over time, and whether they are crowding out public services. I collected up to twelve years of annual financial documents for a diverse set of 442 cities and counties throughout the country, and then I combed each one for information about what they spent on their employees' pensions.

A descriptive analysis of these data reveals a great deal that is new. First, even in the early years covered by the dataset, cities and counties varied tremendously in how much they spent on their employees' pensions, with the top 20-25% spending much larger amounts per employee and as a share of revenue than the rest. Second, it is true that cities and counties in states with strong public-sector unions and mandatory collective bargaining spent more per employee on retirement costs, even as of 2007. Third, city and county pension costs have risen just about everywhere since 2005, but again there is substantial variation in that growth, and the growth has been more rapid in places with strong unions and collective bargaining.

By connecting these local annual pension cost data to U.S. Census data on local government employment and finances, I also find that the average local government is not responding to larger pension cost increases by increasing general revenue. Instead, rising per-

employee pension expenditures are associated with declines in public-sector employment—especially in cities and counties with collective bargaining and high rates of public-sector union membership. In municipal governments, pension-related employment reductions have affected both public safety and non-public safety functions, whereas in counties—which typically have smaller shares of public-safety employees than cities—the reductions have been more pronounced among non-safety employees. Finally, in cities with collective bargaining, greater increases in pension expenditures are also linked to larger reductions in capital outlays.

These findings advance our understanding of the impacts of public-sector pensions in several ways. At a most basic level, these new data enable a discussion of the tradeoffs of public pension costs that is relatively easy to understand. Discussions of public pensions in recent years have focused almost entirely on quantities such as unfunded liabilities and actuarial assumptions—ones that are difficult for citizens, policy practitioners, and government employees to grasp. Most people probably do not have a good sense of what an unfunded pension liability is, why they should care about the discount rate, or what the implications of pension underfunding are likely to be for them. It doesn't help that the consequences of pension underfunding are usually described in terms of events that will happen in the distant future. By comparison, people are probably much more likely to appreciate the significance of how much their governments spend on pensions, how that has changed, and how their local governments are reducing jobs and capital outlays as a result. Putting the focus on what local government is doing less of—not in the future, but right now—in response to pension costs thus has potential to encourage greater public engagement on the issue.

These findings also put the spotlight on important tradeoffs faced by public-sector employees and their unions. Debates about public pensions are often framed as pitting pro-

employee, pro-pension interests against anti-pension, anti-public-worker interests, but my findings here suggest that such a characterization is overly simplistic and misleading. Public-sector employees understandably have an interest in more generous and more reliable retirement benefits, but local governments' payments for those retirement benefits can cut into their ability to offer salary increases and hire more employees (see Anzia and Moe, 2019). My analysis here shows that as local governments spend more on pensions, they have fewer public-sector jobs to offer—an implication that is not positive for government employees or their unions. Thus, by connecting data on local pension costs to data on local employment, we can see clearly the tradeoffs associated with pensions—even for government employees themselves.

A pension-induced transformation of local government as it affects government employees is no small matter, but the implications of rising pension costs are broader than that. Indeed, it stands to affect everyone who relies on local government service provision, including police and fire protection, refuse collection, public parks, libraries, and county court systems. Rising pension expenditures are already changing the landscape of local government, and the findings here suggest that the future of local government may look very different than the past.

References

Andonov, Aleksandar, Yael V. Hochberg, and Joshua D. Rauh. 2018. "Political Representation and Governance: Evidence from the Investment Decisions of Public Pension Funds." *Journal of Finance* 73(5): 2041-86.

Alesina, Alberto, Reza Baqir, and William Easterly. 1999. "Public Goods and Ethnic Divisions." *Quarterly Journal of Economics* 114(4): 1243-84.

Anzia, Sarah F., and Terry M. Moe. 2015. "Public Sector Unions and the Costs of Government." *Journal of Politics* 77 (1): 114-127.

Anzia, Sarah F., and Terry M. Moe. 2016. "Do Politicians Use Policy to Make Politics? The Case of Public-Sector Labor Laws." *American Political Science Review* 110 (4): 763-777.

Anzia, Sarah F., and Terry M. Moe. 2017. "Polarization and Policy: The Politics of Public-Sector Pensions." *Legislative Studies Quarterly* 42 (1): 33-62.

Anzia, Sarah F., and Terry M. Moe. 2019. "Interest Groups on the Inside: The Governance of Public Pension Funds." *Perspectives on Politics*.

De Benedictis-Kessner, Justin, and Christopher Warshaw. 2016. "Mayoral Partisanship and Municipal Fiscal Policy." *Journal of Politics* 78(4): 1124-38.

DiSalvo, Daniel. 2015. *Government Against Itself: Public Union Power and Its Consequences*. New York: Oxford University Press.

DiSalvo, Daniel, and Jeffrey Kucik. 2017. "Unions, Parties, and the Politics of State Government Legacy Cost." *Policy Studies Journal* (November 27).

Gerber, Elisabeth R., and Daniel J. Hopkins. 2011. "When Mayors Matter: Estimating the Impact of Mayoral Partisanship on City Policy." *American Journal of Political Science* 55(2): 326-39.

Hess, David. 2005. "Protecting and Politicizing Public Pension Fund Assets: Empirical Evidence on the Effects of Governance Structures and Practices." *U.C. Davis Law Review* 39: 187-227.

Hoxby, Caroline Minter. 1996. "How Teachers' Unions Affect Education Production." *Quarterly Journal of Economics* 111 (3): 671-718.

Kiewiet, D. Roderick. 2010. "The Day after Tomorrow: The Politics of Public Employee Retirement Benefits." *The California Journal of Public Policy* 2 (3), article 3.

Kiewiet, D. Roderick, and Mathew D. McCubbins. 2014. "State and Local Government Finance: The New Fiscal Ice Age." *Annual Review of Political Science* 17: 105-122.

Lovenheim, Michael F. 2009. "The Effect of Teachers' Unions on Education Production: Evidence from Union Election Certifications in Three Midwestern States." *Journal of Labor Economics* 27(4): 525-87.

Mitchell, Olivia S., and Robert S. Smith. 1994. "Pension Funding in the Public Sector." *The Review of Economics and Statistics* 76 (2): 278-290.

Moe, Terry M. 2011. *Special Interest: Teachers Unions and America's Public Schools*. Washington, DC: Brookings.

Monahan, Amy. B. 2010. "Public Pension Plan Reform: The Legal Framework." *Education Finance and Policy* 5 (4): 617-46.

Munnell, Alicia H. 2012. *State and Local Pensions: What Now?* Washington, DC: Brookings Institution Press.

Munnell, Alicia H., Jean-Pierre Aubry, Josh Hurwitz, and Mark Cafarelli. 2013. "Are City Fiscal Woes Widespread? Are Pensions the Cause?" Center for Retirement Research at Boston College.

Munnell, Alicia H., and Jean-Pierre Aubry. 2016. "An Overview of the Pension/OPEB Landscape." Center for Retirement Research at Boston College, working paper.

Novy-Marx, Robert, and Joshua Rauh. 2009. "The Liabilities and Risks of State-Sponsored Pension Plans." *Journal of Economic Perspectives* 23 (4): 191-210.

Novy-Marx, Robert, and Joshua Rauh. 2011. "Public Pension Promises: How Big Are They and What Are They Worth?" *Journal of Finance* 66 (4): 1207-1245.

Novy-Marx, Robert, and Joshua Rauh. 2014. "The Revenue Demands of Public Employee Pension Promises." *American Economic Journal: Economic Policy* 6 (1): 193-229.

Paglayan, Agustina S. 2019. "Public-Sector Unions and the Size of Government." *American Journal of Political Science* 63(1): 21-36.

Pew Charitable Trusts. 2019. "The State Pension Funding Gap: 2017." Issue brief, June 27, 2019.

Rauh, Joshua. 2016. "Hidden Debt, Hidden Deficits: How Pension Promises are Consuming State and Local Budgets." *Hoover Institution Essay*.

Saltzman, Gregory M. 1985. "Bargaining Laws as a Cause and Consequence of the Growth of Teacher Unionism." *Industrial and Labor Relations Review* 38 (3): 335-351.

Saltzman, Gregory M. 1988. "Public Sector Bargaining Laws Really Matter: Evidence from Ohio and Illinois." In *When Public Sector Workers Unionize*, Richard B. Freeman and Casey Ichniowski, eds. Chicago: University of Chicago Press.

Stalebrink, Odd J. 2014. “Public Pension Funds and Assumed Rates of Return: An Empirical Examination of Public Sector Defined Benefit Pension Plans.” *The American Review of Public Administration* 44 (1): 92-111.

Tausanovitch, Chris, and Christopher Warshaw. 2014. “Representation in Municipal Government.” *American Political Science Review* 108(3): 605-41.

Thom, Michael. 2013. “All of the Above: How Fiscal, Political, and Workforce Traits Affect Pension Funding.” *State and Local Government Review* 45 (3): 163-171.

Trounstine, Jessica. 2018. *Segregation by Design: Local Politics and Inequality in American Cities*. New York: Cambridge University Press.

Walsh, Mary Williams. 2011. “The Little State With a Big Mess.” *New York Times*, October 23, 2011.

Table 1: City and county pension expenditures in 2007

	Pension costs, proportion of general revenue	Pension costs per full-time equivalent employee (2016 dollars)
1%	0.003	$308
5%	0.006	$609
25%	0.018	$2,619
50%	0.031	$4,901
75%	0.047	$7,726
95%	0.077	$15,183
99%	0.109	$21,728
Mean	0.035	$5,945
SD	0.026	$4,996
N	410	409

Table 2: City and county pension expenditures as of 2007

	Proportion of general revenue	Per full-time equivalent employee (thousands of 2016 dollars)	
	(1)	(2)	(3)
Union membership	0.0004	0.548	
	(0.010)	(0.259)	
Collective bargaining			0.194
			(0.106)
Ln(income per capita)	0.013	0.444	0.441
	(0.014)	(0.236)	(0.245)
Democratic presidential vote	0.005	-0.2	-0.21
	(0.012)	(0.335)	(0.351)
% DC plans	-0.02	-0.24	-0.269
	(0.006)	(0.163)	(0.181)
Number of plans	0.001	0.004	0.007
	(0.001)	(0.027)	(0.027)
% Local plan	0.007	0.097	0.11
	(0.004)	(0.093)	(0.094)
% also funding OPEB	0.001	0.119	0.137
	(0.006)	(0.101)	(0.100)
Ln(population)	0.0001	0.086	0.095
	(0.002)	(0.036)	(0.037)
% Urban	0.011	0.351	0.35
	(0.006)	(0.207)	(0.208)
% Homeowner	0.014	0.067	0.037
	(0.015)	(0.386)	(0.405)
% Black	-0.008	-0.19	-0.255
	(0.011)	(0.286)	(0.300)
% Asian	0.013	0.802	1.035
	(0.025)	(0.433)	(0.443)
% Hispanic	0.027	0.729	0.638
	(0.020)	(0.384)	(0.381)
County	-0.013	-0.48	-0.495
	(0.004)	(0.104)	(0.104)
Constant	-0.121	-4.137	-4.088
	(0.132)	(2.147)	(2.207)
R-squared	0.26	0.49	0.48
Observations	407	406	406

Notes: Standard errors clustered by state in parentheses.

Table 3: Within-city and county change in pension expenditures

	Proportion of general revenue, 2005-2016	Per employee, 2005-2016 (thousands)		Per employee, 2007-2016 (thousands)
	(1)	(2)	(3)	(4)
Union membership	0.019	3.953		3.002
	(0.009)	(1.345)		(1.066)
Collective bargaining			0.884	
			(0.570)	
Ln(income per capita)	0.001	1.017	0.783	0.006
	(0.005)	(1.218)	(1.245)	(0.860)
Democratic presidential vote	-0.011	-0.546	0.579	-0.086
	(0.012)	(2.024)	(2.425)	(1.877)
% DC plan	-0.012	-3.004	-3.308	-2.848
	(0.006)	(0.999)	(1.039)	(0.930)
Number of plans	0.001	0.134	0.15	0.079
	(0.001)	(0.125)	(0.126)	(0.107)
% Local plans	0.013	1.689	1.835	2.469
	(0.005)	(0.653)	(0.656)	(0.549)
% also funding OPEB	0.014	1.241	1.441	1.521
	(0.007)	(0.916)	(0.849)	(1.264)
Ln(population)	-0.001	0.085	0.17	-0.073
	(0.002)	(0.271)	(0.264)	(0.213)
% Urban	-0.003	-0.543	-0.613	-0.944
	(0.006)	(1.136)	(1.097)	(0.873)
% Homeowner	0.012	-0.447	-0.143	0.335
	(0.017)	(2.756)	(2.791)	(2.345)
% Black	0.025	-0.023	-1.318	1.86
	(0.014)	(1.643)	(1.768)	(1.744)
% Asian	0.039	13.746	15.664	8.534
	(0.017)	(3.499)	(3.329)	(4.193)
% Hispanic	0.024	5.977	5.092	4.709
	(0.016)	(2.745)	(2.993)	(1.887)
County	-0.002	-0.836	-0.855	-1.098
	(0.004)	(0.666)	(0.677)	(0.508)
Constant	-0.01	-11.006	-9.24	1.101
	(0.043)	(10.276)	(10.797)	(7.428)
R-squared	0.21	0.42	0.4	0.34
Observations	348	265	265	351

Notes: Standard errors clustered by state in parentheses.

Table 4: Pension costs and general revenue

	Ln(General revenue per capita) (1)	Ln(Own-source general revenue per capita) (2)	Ln(General revenue per capita) (3)	Ln(Own-source general revenue per capita) (4)
Ln(Pension expenditures per employee)	-0.001	0.008	0.028	0.009
	(0.022)	(0.019)	(0.056)	(0.051)
Collective bargaining			-0.037	-0.001
* Ln(Pension exp. per employee)			(0.062)	(0.054)
Ln(income per capita)	0.72	0.706	0.716	0.706
	(0.196)	(0.187)	(0.194)	(0.187)
Ln(population)	-0.683	-0.722	-0.686	-0.722
	(0.140)	(0.148)	(0.138)	(0.148)
% Urban	0.183	0.435	0.16	0.434
	(0.402)	(0.400)	(0.406)	(0.403)
% Homeowner	0.056	0.002	0.061	0.003
	(0.233)	(0.274)	(0.233)	(0.273)
% Black	1.083	1.168	1.072	1.167
	(0.790)	(0.746)	(0.779)	(0.744)
% Asian	0.323	-0.037	0.347	-0.036
	(0.521)	(0.501)	(0.523)	(0.505)
% Hispanic	0.339	0.199	0.364	0.2
	(0.370)	(0.449)	(0.376)	(0.451)
R-squared	0.98	0.97	0.98	0.97
Observations	4,133	4,133	4,133	4,133

Notes: Standard errors clustered by state in parentheses. Models include local government and year fixed effects.

Table 5: Pension costs and local government employment

	Full-time equivalent (1)	Full-time (2)	Part-time (3)	Full-time (4)	Full-time (5)
Ln(Pension expenditures per employee)	-0.066	-0.068	-0.009	-0.029	-0.009
	(0.014)	(0.013)	(0.039)	(0.023)	(0.027)
Collective bargaining				-0.051	
* Ln(Pension exp. per employee)				(0.026)	
Union membership					-0.118
* Ln(Pension exp. per employee)					(0.057)
Ln(income per capita)	0.328	0.355	0.074	0.351	0.349
	(0.091)	(0.105)	(0.258)	(0.105)	(0.103)
Ln(population)	-0.634	-0.63	-0.824	-0.633	-0.634
	(0.105)	(0.107)	(0.177)	(0.103)	(0.101)
% Urban	0.431	0.472	0.949	0.44	0.437
	(0.172)	(0.179)	(0.791)	(0.176)	(0.173)
% Homeowner	0.012	0.021	0.217	0.027	0.04
	(0.137)	(0.144)	(0.451)	(0.145)	(0.146)
% Black	0.687	0.623	2.22	0.609	0.606
	(0.388)	(0.353)	(1.293)	(0.350)	(0.348)
% Asian	0.081	-0.024	-0.32	0.012	0.024
	(0.576)	(0.575)	(1.080)	(0.560)	(0.550)
% Hispanic	-0.064	-0.16	0.542	-0.119	-0.117
	(0.338)	(0.342)	(1.085)	(0.339)	(0.346)
R-squared	0.99	0.99	0.94	0.99	0.99
Observations	4,075	4,075	4,041	4,075	4,075

Notes: Standard errors clustered by state in parentheses. Models include local government and year fixed effects.

Table 6: Pension costs and local government employment, by local government and employee type

	Municipal police protection	Municipal fire protection	Municipal non-public safety	County police protection	County corrections	County non-public safety
	(1)	(2)	(3)	(4)	(5)	(6)
Ln(Pension expenditures per employee)	0.033	-0.006	-0.01	0.107	-0.322	-0.047
	(0.028)	(0.051)	(0.021)	(0.088)	(0.148)	(0.049)
Collective bargaining	-0.082	-0.065	-0.086	-0.092	0.271	-0.016
* Ln(Pension exp. per employee)	(0.032)	(0.054)	(0.029)	(0.097)	(0.166)	(0.054)
Ln(income per capita)	0.05	0.475	0.175	0.215	0.115	0.425
	(0.147)	(0.255)	(0.121)	(0.192)	(0.274)	(0.178)
Ln(population)	-0.697	-0.652	-0.657	-0.256	-0.872	-0.449
	(0.090)	(0.122)	(0.093)	(0.389)	(0.246)	(0.217)
% Urban	1.468	-2.081	0.357	0.024	-0.011	-0.014
	(0.398)	(0.540)	(0.347)	(0.827)	(1.435)	(0.318)
% Homeowner	0.066	-0.481	0.025	-0.246	0.779	0.308
	(0.175)	(0.412)	(0.204)	(0.490)	(0.722)	(0.237)
% Black	0.27	0.788	0.11	4.128	-0.51	2.216
	(0.426)	(0.540)	(0.280)	(2.658)	(1.957)	(1.198)
% Asian	0.633	0.824	0.25	3.79	-2.172	-1.255
	(0.482)	(0.669)	(0.547)	(1.802)	(1.712)	(0.773)
% Hispanic	0.245	-0.364	-0.269	0.682	-0.053	0.292
	(0.349)	(0.626)	(0.408)	(1.077)	(1.056)	(0.507)
R-squared	0.97	0.98	0.99	0.93	0.86	0.99
Observations	1,876	1,675	1,889	2,124	1,975	2,186

Notes: Standard errors clustered by state in parentheses. Models include local government and year fixed effects.

Table 7: Pension costs and capital outlays

	(1)	(2)
Ln(Pension expenditures per employee)	-0.163	0.308
	(0.224)	(0.535)
Collective bargaining	-0.138	-0.098
* Ln(Pension exp. per employee)	(0.255)	(0.533)
Ln(income per capita)	2.269	1.311
	(0.609)	(1.081)
Ln(population)	-0.987	-1.693
	(0.548)	(0.675)
% Urban	-3.227	-1.41
	(2.420)	(1.694)
% Homeowner	-0.125	0.74
	(1.190)	(2.655)
% Black	4.352	-1.21
	(3.299)	(6.891)
% Asian	-3.673	5.4
	(2.467)	(5.708)
% Hispanic	3.826	-2.49
	(2.414)	(2.963)
R-squared	0.63	0.69
Observations	1,930	2,136

Notes: Standard errors clustered by state in parentheses. Models include fixed effects for local government and year.

Figure 1: Within-city and county change in pension expenditures