

**Pensions in the Trenches:
Are Rising City Pension Costs Crowding Out Public Services?**

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

August 6, 2017

Abstract: Some experts claim that state and local governments have seen dramatic increases in their public pension costs and that pension spending is crowding out local public services. Others maintain that serious pension problems are limited to a small number of governments that have been especially irresponsible. However, no existing studies—nor the datasets they rely on—allow us to evaluate the extent to which local pension costs are actually rising, or whether pensions are crowding out services. In this paper, I analyze a new dataset of the annual pension expenditures of 219 municipal governments across the U.S. from 2005 to 2014. I find that 85% of the cities saw increases in their pension expenditures over this ten-year period. In the median city, inflation-adjusted pension expenditures increased by 45% in ten years, and the average increase was 69%. There is also considerable cross-city variation in the amount cities spend on pensions per employee, with typical cities spending about \$7,000 per employee per year but some spending as little as \$2,500 and others spending more than \$20,000. And when I examine variation within cities over time, I estimate that a 10% increase in per-employee pension expenditures is associated with a 0.73% average drop in city employment the following year. These pension-induced employment reductions are most pronounced for non-public safety employees and for cities in states with collective bargaining laws. In addition, many cities are seeing cuts in areas other than employment: rising pension costs are also associated with reduced spending on construction and equipment.

This paper was prepared for the annual conference of the American Political Science Association, August 31-September 3, 2017. Many thanks to the Berkeley Institute for the Future of Young Americans and the Institute for Research on Labor and Employment for providing funding for this project.

Many Americans take for granted that their local governments will provide public services like police protection, fire protection, street sweeping, and refuse collection. But in the years to come, it may well become harder for local governments to carry out these basic functions. One of the main reasons, according to scholars, journalists, and good government groups, is the substantial and growing cost of government employees' retirement benefits.

Unlike in the private sector, where defined contribution plans are most common, nearly all employees of state and local governments have traditional pensions, or defined benefit plans. And over the last few decades, state and local officials have enacted changes to make pension benefits more generous (Anzia and Moe 2017a). These changes have not only made pensions more expensive for state and local governments, but they have also proven extremely difficult—if not impossible—to roll back. In most states, public pensions are backed by strong legal guarantees, such that benefit formulas can only be changed for future hires (Monahan 2010). Once governments adopt more generous benefits, they are locked in to paying for them.

What is more, governments across the country have consistently underfunded their pensions, setting aside far too little money to pay for the benefits they have promised. A variety of decisions have contributed to the underfunding, including strategically-chosen actuarial assumptions that make pension liabilities look smaller than they actually are (e.g., Novy-Marx and Rauh 2009, 2011), failure to pay the annual amounts required for full funding (e.g., Anzia and Moe 2017b), and possibly politically-motivated investment decisions (e.g., Andonov et al. 2016). The upshot is that in total, public employee pensions in the United States are now underfunded by somewhere between \$934 billion and \$3.4 trillion, depending on the estimates used (e.g., Rauh 2016, Pew Charitable Trusts 2016). Because the benefits are legally guaranteed,

someone has to make up for the shortfalls—and many experts claim that state and local governments are shouldering the burden in the form of large and rising pension expenditures.

Based on a few cities and states that have received a great deal of media attention, it is clear that in some places pension costs are substantial and have risen rapidly—and that certain governments have struggled to cope with the fiscal pressure. Cities like San Jose and San Diego have tried to reform pensions by reducing benefits for current employees and raising employee contribution rates—moves that landed those governments in a quagmire of legal challenges.¹ Chicago recently approved the largest property tax increase in its history, and Oakland has laid off police officers and issued pension obligation bonds. Detroit, Vallejo, and Stockton have declared bankruptcy. And in all of these cities, skyrocketing pension costs have played a major role in forcing government officials to make painful decisions about pension benefits, overall service provision, government employment levels, taxes, and debt.

By some accounts, the events unfolding in places like San Jose, San Diego, and Chicago are part of a much larger—and permanent—trend. Kiewiet and McCubbins (2014) argue that the aging of the population and decades of irresponsible decision-making have led to dramatic increases in governments' pension costs. What's more, they argue, that trend will only continue:

“To describe what is happening in state and local government finance as a crisis is...something of a misnomer. A crisis is an acute episode, a period of travail that must somehow be endured until normal financial conditions are restored...What we are experiencing is the onset of the New Fiscal Ice Age, a period in which a given level of state and local tax revenue purchases a considerably lower level of current services...In most places, fiscal conditions will become increasingly harsh” (Kiewiet and McCubbins 2014, p. 106).

¹ Most governments only change the benefits of new hires, which does little to address short-term fiscal problems.

This sentiment is echoed by other experts on public pensions. The nonpartisan Little Hoover Commission of California has warned that “pension costs will crush government: government budgets are being cut while pension costs continue to rise and squeeze other government priorities.”² A 2011 *New York Times* article about Rhode Island explains that because of rising pension costs, “many Americans may be forced to rethink what government means at the state and local level.”³ And in an account of how pension and health care costs are putting pressure on government budgets, DiSalvo (2015) raises examples from governments ranging from Oakland to Chicago to New York to Scranton, Pennsylvania.

Others argue that these claims of widespread pension-induced fiscal crisis are exaggerated. Munnell et al. (2013) point out that we should not draw general conclusions about the state of pension costs in the United States after only examining what may well be the worst cases. Based on their own study of municipal fiscal stress, Munnell et al. (2013, p. 5) conclude, “The question is whether cities across the country are about to topple like dominoes. And whether pensions are the problem. The answer appears to be ‘no’ on both fronts.” That assessment echoes that of Munnell’s book on state and local pension finances (2012): she concludes that while the states and cities with the worst-managed pension funds dominate the news cycle, most state and large local pension funds are actually in fairly good condition.

Which of these conclusions is closer to the mark? Are local governments across the country entering a New Fiscal Ice Age—a bleak environment in which most governments are struggling to cope with pensions costs? Or are the serious fiscal challenges confined to a few states and large cities—and caused by problems other than pensions? Given the state of existing

² Little Hoover Commission, “Public Pensions for Retirement Security,” February 2011, p. iii.

³ Walsh, Mary Williams. 2011. “The Little State With a Big Mess.” *New York Times*, October 23, 2011.

research, it is impossible to know. While there is a large literature on public employee pensions, none of its studies—nor the datasets they rely on—allow evaluation of how much cities are spending on pensions, how pension expenditures vary across cities, the degree to which those expenditures are rising, or whether pension spending is crowding out public services.

In this paper, I begin to answer those questions by analyzing a new dataset of the annual pension expenditures of 219 municipal governments from 2005 to 2014, which I hand-collected from the cities' annual financial statements. This dataset is unlike any that existed before, because it tracks actual city 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 across the country. With these new data, we can see for the first time what American cities are spending on their employees' pensions and how that has changed over a 10-year period.

In the first empirical section of the paper, I carry out a descriptive analysis of city pension costs over time. I find that 85% of the cities in the sample saw increases in their pension expenditures over this ten-year period. In the median city, inflation-adjusted pension expenditures increased by 45% in ten years, and in the average city, the increase was 69%. I also find tremendous cross-city variation in the levels of pension expenditures. While median annual pension spending for cities in this sample is about \$7,000 per employee per year, some cities spend as little as \$2,500, while a few spend more than \$20,000. And in my analysis of variation in per-employee pension expenditures across cities, I find that larger, more urban areas spend more on pensions for each of their employees, as do cities in states with collective bargaining laws and strong public-sector unions.

In the next part of the empirical analysis, I examine variation within cities over time to evaluate the extent to which rising pension costs are crowding out public services. I find that in

the average city, increases in per-employee pension expenditures are significantly associated with decreases in city employment the following year. But the effects of rising pension expenditures are not uniform across cities or across types of city employees. The negative effect of rising pension costs on city employment is significantly greater for cities in states with collective bargaining laws. Also, the employment reductions typically occur more so among non-safety employees than among fire and police protection employees. Moreover, there are visible effects of rising pension costs on city construction and equipment spending: for cities in collective bargaining states, larger increases in pension expenditures are associated with bigger decreases in capital outlays. Together, these findings show that rising pension costs are indeed crowding out public services in many places—but that the average effects are more pronounced for certain cities and certain kinds of employees.

Literature

While there is a sizeable literature on public employee pensions, most existing research examines the funding ratios and actuarial assumptions of state-administered pension plans and the largest locally-administered plans. There are good reasons for this. The vast majority of state and local government employees in the United States are enrolled in these large plans, such as CalPERS in California and OPERS in Ohio. And even by the plans' own estimates, most state-operated plans do not have sufficient assets to cover the retirement benefits that have been promised. Explaining variation in plans' fiscal health is therefore an important subject of research, and several scholars have attempted to identify the political, economic, and governance factors that influence the degree of plans' underfunding (e.g., Thom 2013, Stalebrink 2014, Andonov et al. 2016, Mitchell and Smith 1994).

One topic of considerable debate in this literature has to do with what public pension liabilities are worth. Many finance experts argue that the plans' actuarial assumptions make their liabilities look far smaller than they actually are—thereby making their funding ratios look better than they actually are. Research by Novy-Marx and Rauh (2009, 2011, 2014), for example, develops revised estimates of pension liabilities and funding ratios, using more appropriate assumptions, and it also projects how much government pension contributions would have to increase in order for them to work toward full funding. By the account of Novy-Marx and Rauh (2014), pension contributions would have to grow by 2.5 times in order for state and local pension systems to achieve full funding over 30 years. Because these projections are so different than the ones governments themselves report—and because the implications of the revised projections are so dire—this research has led to intense debate about how to properly project pension liabilities and report funding ratios.

An equally contentious issue in the literature centers on the role of government employees, public-sector unions, and collective bargaining in influencing pension benefits and funding levels. On the one hand, research has shown that collective bargaining and strong public-sector unions lead to higher salaries and especially higher fringe benefits for government employees (e.g., Anzia and Moe 2015). This research does not examine pension benefits directly because of data limitations, but some scholars have argued that strong unions should naturally have the same effect on pension benefits (e.g., Kiewiet 2010). In a finding consistent with that hypothesis, Anzia and Moe (2017a) show that public-sector unions influence state legislators' votes on pension benefit changes—specifically by making them less favorable toward benefit reductions. But the issue is far from settled. Munnell (2012) finds that state collective bargaining laws are not significantly associated with a higher “normal cost” of pensions as a

share of payroll (a measure of pension benefit generosity).⁴ There is also disagreement about how public-sector unions affect pension *funding*. The dominant claim in the literature is that government employees want to ensure that their pensions are well-funded—and thus that strong unions should be associated with better-funded pensions (Hess 2005, Romano 1995). Anzia and Moe (2017b) question this, however, arguing that government employees actually have strong incentives to underfund their own pensions. In a study of the funding decisions made by state pension boards, they find that strong public-sector unions and greater government employee representation on the boards are associated with *less* responsible funding decisions.

This discussion makes it clear that scholars are engaged in productive debate about the magnitude and political economy of the nation’s pension problems. But existing research is overwhelmingly focused on the extent of and variation in plan underfunding, and so we still have little sense of what local governments are experiencing and how they are reacting to any changes in pension costs. The most likely reason for the shortage of work in this area is that there aren’t any data on pension costs at the level of the local government. In fact, nearly all of the quantitative empirical work on public pensions relies on the Public Plans Database developed by Boston College’s Center for Retirement Research, which compiles key statistics from the Comprehensive Annual Financial Reports (CAFRs) of the major state- and locally-operated pension plans. This is a rich dataset in that it documents each plan’s funded ratios, actuarial assumptions, required contributions, and more—for all state and major local plans back to the early 2000s. But these plan-level data do not tell us about the pension expenditures of particular governments. It is of course reasonable to assume that if pension benefits have expanded and unfunded liabilities have increased, then many local governments are probably being asked or

⁴ Normal cost is the value of pension benefits earned by employees in a given year.

required to make larger annual contributions than in the past. But we don't know how pervasive the pattern has been, how large the increases are, how governments are responding to any such changes—or, for that matter, how much governments even spend on pensions.

As I discussed earlier, some scholars argue that pension costs are a heavy burden for local governments, that increases in pension costs are significant and widespread, and that cities, counties, and special districts throughout the country are facing difficult budgetary choices as a result. But these studies either rely on state- or plan-level data or examples from a few cities to draw that conclusion (Kiewiet and McCubbins 2014, DiSalvo 2015). And as Munnell et al. (2013) point out, the experiences of cities like New York might not be typical of other cities. Moreover, while state- and plan-level data are indeed helpful for documenting that pension costs are rising in many states, they do not directly speak to what individual local governments are experiencing, how their experiences vary, and how they are coping.

It is worth emphasizing that the problem is not a general lack of data on public pensions; rather, it is a mismatch between the unit of analysis in available datasets—which is the pension *plan*—and the unit of analysis that is needed to answer questions about local government fiscal stress—which is the individual local *government*. State-operated pension plans like CalPERS and OPERS have many local governments as participants, and it is usually not possible to accurately attribute parts of these plans' liabilities, assets, and contributions to their hundreds of local government participants. What is more, many individual local governments contribute to multiple plans for different categories of employees, some of which are state plans, others of which are local plans. In order to say anything definitive about pension costs at the level of the individual local government—which is what we need if we want to compare the costs of different governments and explore changes in their costs over time—we would not only need to

apportion state plans' contributions to each of their participating local government members, but we would also need to do that for all state plans in which a local government participates, as well as for any local plans to which it contributes.

In a recent project, Munnell and Aubry (2016) pursue this kind of strategy in an attempt to estimate what pension costs look like at the level of individual governments. For all 50 states and a sample of 174 cities, 178 counties, and 415 school districts, they estimate the share of state plans' annual required contributions in the year 2014 that were the responsibility of each individual government. Then, applying a uniform discount rate and amortization period to all plans, they estimate what each government in their sample would have had to have paid in 2014 in order to fully fund their pension obligations.

This approach has many strengths, first that it uses existing data to come up with estimates of pension costs at the local level, and second that it results in measures of what individual governments *should* be paying—not just what they *are* paying, which is often too little. However, to the extent that Munnell and Aubry are focusing on “major” cities (which are presumably large cities), their findings may not generalize to smaller local governments, and it is also unclear whether their estimates include costs of locally-administered plans, which are common especially in the largest cities and counties. Most importantly, because they carry out this complicated estimation for a single year—2014—the data are not suited to answering questions about change over time: one cannot use their estimates to determine whether, and the extent to which, pension costs have *risen* over the last several years.

It is also worth drawing a distinction between what governments *should* be paying and what they *are* paying—and we should ask which is the more appropriate quantity for answering questions about what has been happening to local government budgets across the country. It is

readily apparent that many governments have not been contributing what they should—hence the severe underfunding of many plans. And for answering some questions, it makes sense to focus on what individual governments *should* be contributing, as Munnell and Aubry (2016) have done. But here I am asking whether pension expenditures have risen over time and whether pension spending is crowding out public services. Thus, the more appropriate measure is what cities are *actually* spending on pensions, because it is the expenditures that would crowd out services. So to answer *these* questions, we not only need data at the level of the individual government, but more precisely we need data on individual governments’ *actual* pension expenditures over a period of time. Those, after all, are the numbers affecting their budgets.

In sum, the existing literature pursues certain pension-related questions: questions about the value of liabilities and the degree of underfunding, and questions about the political and economic factors that have contributed to underfunding and benefit levels. Yet because of data limitations, we know very little about how pension expenditures vary from one local government to another, nor has anyone systematically assessed the extent to which pension costs have risen. Some claim that local governments are being forced to spend larger and larger sums of money each year on their employees’ pensions—and that that has forced draconian cuts to services (e.g., Kiewiet and McCubbins 2014, DiSalvo 2015). Others argue that such claims about pension-induced fiscal stress are exaggerated, and that many or even most local governments are not experiencing such stress (e.g., Munnell et al. 2013, Munnell and Aubry 2016). Without any data on what cities spend on pensions, we cannot say for sure which is closer to the mark.

Data

To evaluate how local governments’ pension costs have changed over time, and how any such changes have affected city services, I assembled a new dataset. My goal for assembling the

dataset was to collect several recent years of the pension expenditures of a diverse set of municipal governments in the United States. Unfortunately, there is no central repository for such information, so I set out to collect a large number of cities' annual financial statements—usually called CAFRs—which detail what the cities spent on their various employee pension plans in each year.

City CAFRs are the only reliable source of information on cities' pension expenditures, but it can be difficult to locate them, and sometimes costly to acquire them. Once the CAFRs are in hand, moreover, it takes time to locate the relevant information and interpret it, because most CAFRs are hundreds of pages long. Collecting and reading the CAFRs of thousands of municipal governments for several years would therefore have been prohibitively costly. At the same time, an important goal of this project is to determine how common it is for cities to have rising pension costs, and so I did not want to limit the data collection to a small number of cities.

With these tradeoffs in mind, I selected a sample of 236 municipal governments from the 1,915 that appear in the U.S. Census Bureau's Survey of Governments Finance files for every year between 2005 to 2012.⁵ I first defined eight strata based on city population, with the first stratum being cities with fewer than 10,000 residents and the last being cities with more than 1 million residents. I then used random sampling with replacement to draw 40 municipal

⁵ The U.S. Census collects finance and employment information for all governments during its census every five years, but it also collects the information from a sample of governments during years in between censuses. For its annual Finance survey in recent years, the U.S. Census used one consistent sample of local governments for the 2005-2008 time period and a different sample for the 2009-2012 time period. (At the time I selected the sample of cities for CAFR collection, 2012 was the latest year available.) I therefore determined which municipal governments were included in both of these samples and then drew 236 municipal governments from that set.

governments from each stratum, weighting by population within strata. This produced a sample of 236 cities from all 50 states as well as Washington, DC.⁶

For each of those 236 municipal governments, I set out to collect the CAFRs for each year from 2005 to 2014. Most cities had at least some CAFRs available on their websites, typically for the two or three most recent years, but some cities' websites did not provide any CAFRs at all. For city-years for which the CAFRs were not available online, I contacted the cities to request the documents, filing public information requests where necessary. Many cities provided their CAFRs at no charge; others provided them for a small fee; others requested substantial fees; and still others did not respond to my requests. In the end, I was able to obtain at least some years' CAFRs for 219 of the municipal governments in my sample (93%). For 190 municipal governments I acquired the full set of CAFRs for all ten years; for the remaining 29, I obtained CAFRs for some years (usually the later years) but not others.

In a typical CAFR, the information about city pension expenditures is available in the Notes to the Financial Statements. While different CAFRs present the information in different formats—with some providing tables and others describing pension costs in a narrative—most CAFRs provide certain key pieces of information: the retirement plans to which the city contributed on behalf of its employees (most cities have more than one, with some cities having up to 10 plans), some background on plan history and eligibility criteria, each plan's Annual Pension Cost or Actuarially Required Contribution, and the amount of the employer contribution to each plan in that year. Therefore, the main piece of information I drew from the CAFRs was the amount the city contributed to each of its employee retirement plans in that year.

⁶ Some cities (especially large cities) were selected more than once, which is why the number of unique cities in the sample (236) is less than 320.

A few caveats are worth highlighting. The first is that my intention was to collect cities' *total* pension contributions, including any amount of the city's employee contributions paid by the city (typically called Employer-Paid Member Contributions, or EPMC). However, I discovered that in most CAFRs, it is difficult to discern whether the city is picking up any of the employees' share of contributions, and even when a CAFR does indicate the presence of EPMC, the dollar value of any EPMC is typically not reported clearly and consistently. Most likely, getting reliable information about the prevalence and cost of EPMC would require access to cities' collective bargaining agreements, which is beyond the scope of this study. For practical reasons, therefore, the retirement contributions discussed below mostly do not include EPMC.

In collecting the data, I had to make other decisions about what to include or exclude. For example, a few states, such as Ohio, fund their other post-employment benefits (OPEB)—such as retiree healthcare—using money cities contribute to their pension funds. My focus here is on pensions and related retirement benefits, not on OPEB. Therefore, wherever possible, I subtracted any funds going to OPEB from city pension contributions, as long as I could consistently do that for all years within a city-plan. However, for a few city-plans I was not able to back out the funds going to OPEB, and so there are a small number of cases in which OPEB expenditures are included. In addition, I had to decide whether to include city contributions to defined contribution (DC) plans. While DC plans in these cities are rare and typically make up a very small share of city retirement contributions, they are a potential alternative to defined benefit plans, and so I include any city DC plan contributions in the dataset.

The final challenge in this data collection effort relates to pension obligation bonds. These bonds, commonly referred to as POBs, are taxable bonds issued by governments to pay their obligations to their pension plans. POBs are fairly common across the U.S., but their use is

highly controversial, and increasingly, experts view them as fiscally irresponsible (e.g., Kiewiet 2010). When a city issues POBs, it typically makes a very large one-time payment to its pension fund using the bond revenues, with the hope being that the investment returns on those revenues will exceed the cost of servicing the debt. This sort of gamble has backfired for cities like Oakland and Detroit, in effect leading to ever-higher required pension contributions in the future.

While it is not my goal here to assess the prudence of POBs, POB use by cities in my dataset does require careful consideration. Most importantly, cities that issue POBs typically use the revenue to make a large non-recurring pension contribution in that year, and so in terms of their pension expenditures, it looks like they are suddenly spending much more on pensions. Also, if the ultimate goal is to track cities' total pension costs in each year, it would make sense to include the costs of servicing that debt. In an attempt to collect information on pension-related debt in each city-year, therefore, I also consulted the Long-Term Debt section of each CAFR.⁷ From that part of the CAFR, I was able to track whether a city has pension-related debt, when the bonds were issued, and the amount of the bonds. However, there was no way to consistently track city expenditures on debt service: most CAFRs do not break out POB servicing payments into principal and interest, and because revenue from POBs is included in my pension expenditure data, including payments of the principal would amount to double-counting. Therefore, my dataset does not include city interest payments on pension-related debt.

⁷ Typically pension bonds are called pension obligation bonds, but I included any form of bond issue used to fund pension benefits. For example, in Pembroke Pines, Florida, the city issued bonds called "Taxable Communications Services Tax Revenue Bonds," but the proceeds were used to fund pensions. If some cities issued bonds for multiple purposes—one of which is funding pensions—I was not able to identify that as pension-related debt.

Even with these caveats, which are unavoidable given the complexity of city pension expenditures,⁸ this dataset is unlike any that existed before, and it is uniquely suited to the task of assessing the experiences of American cities. In total, I have a dataset of 2,081 annual pension expenditure observations from 219 unique cities. With these data, I will be able to see what cities are spending on their employees' pensions and how that changed over this ten-year period.

City Pension Expenditures, 2005-2014

What, then, do these data tell us about pension expenditures in these 219 cities? As a first step, for purposes of illustration, Figure 1 presents the annual pension expenditures of 16 cities of different sizes, with all pension expenditures adjusted to 2014 dollars. (The scale of the vertical axis is different for each graph.) The top row shows the over-time patterns in four cities of fewer than 25,000 residents: Ulysses, Kansas (population 6,161), Cody, Wyoming (population 9,520), Eastlake, Ohio (population 18,577), and Coralville, Iowa (population 18,907). In three out of the four cities, pension expenditures have risen over this ten-year period. Ulysses, Kansas, has seen its pension expenditures increase from \$47,190 in 2005 to \$166,000 in 2014—a 250% increase. In Cody, Wyoming, expenditures have risen 55%, from \$511,290 in 2006 to \$793,000 in 2014.⁹ And in Coralville, Iowa, pension expenditures nearly doubled from \$453,750 in 2005 to \$905,000 in 2014. Eastlake, Ohio, is the only city in this group that has seen little change in its pension expenditures over time.

In row 2 of Figure 1, I find similar patterns for four cities with between 25,000 and 100,000 residents. Pension expenditures in Big Spring, Texas, jumped from \$1,076,665 in 2005 to \$2,462,713 in 2014. In Lacey, Washington, pension expenditures *quadrupled* between 2006

⁸ Also, all of these caveats—with the exception of the occasional city whose pension expenditures include OPEB—work to understate cities' pension expenditures.

⁹ I am missing data on 2005 pension expenditures in Cody, Wyoming.

and 2014.¹⁰ A mid-sized city in Illinois—Palatine—saw its pension expenditures grow from \$3.6 million to \$7.4 million in ten years. And in Warwick, Rhode Island, costs increased but by a more modest 43% from 2005 to 2014.

What of larger cities, such as Anchorage, Alaska, or Miami, Florida? The third row of Figure 1 shows that pension expenditures are rising in cities of this size as well. Peoria, Arizona, a city with 154,000 people, has seen a steady increase in its pension costs, as has Anchorage, Alaska. In Buffalo, New York, pension expenditures shot up suddenly in 2012 and 2013. The only exception here is Miami, Florida, which has paid less toward its pensions since 2011 than it did before the Great Recession.

Finally, in the bottom row of Figure 1, I present the over-time patterns for four of the largest cities in the United States. The patterns in Seattle and San Jose mirror those of the smaller cities: San Jose spent \$102 million (in 2014 dollars) on pensions in 2005 and \$231 million in 2014. Seattle's expenditures increased from \$71 million to roughly \$120 million in the same ten-year period. But the figures from Dallas and Chicago look different—and it is worth discussing how to interpret them. In Dallas, the city issued POBs in 2005, which explains its very large pension expenditure that year—which makes it look as though Dallas is spending less on its pensions today than it did in 2005. Chicago is another interesting example, because it has attracted national media attention for its fiscal woes—largely attributable to pension costs—and yet has been spending less on pensions in recent years than in the past. But this is the caveat that needs to be kept in mind with this dataset: some cities, especially those in the most serious trouble, pay less and less toward their pensions as their fiscal problems worsen. Thus, in cities

¹⁰ Here, too, I am missing data on the city's 2005 pension expenditures.

like Chicago, there is almost certainly a very large gap between what the city is actually paying (which is what I have collected) and what it should be paying.

It is therefore not appropriate to use these data to draw conclusions about which cities are “in trouble” or not. Eastlake, Ohio, is another example: part of the reason that Eastlake’s pension expenditures have not changed as much as other cities’ is that Ohio has statutes setting out local governments’ contribution rates to state-operated pension plans—in the form of a fixed percentage of payroll. The upshot may be that its pension plans are more severely underfunded than plans in other states. But again, for the questions I am asking here—about what is happening with local government budgets—the focus should be on what cities are actually paying. And as it turns out, patterns like the one in Chicago appear to be unusual. In other large cities, such as in San Francisco and New York, pension expenditures are rising steadily.

From Figure 1, therefore, it is clear that cities of different sizes across the country are experiencing increases in their pension expenditures. Turning now to the full sample, what can we say more generally about how city pension expenditures have changed within cities over time? To answer this question, I calculate for each city the percentage change in its pension expenditures from the beginning to the end of the ten-year time period. Because a few cities make large non-recurring pension expenditures in a single year, I first create a measure of each city’s “typical” expenditures in the early period as well as a measure of each city’s “typical” expenditures in the later period. For the former, I calculate the average annual city pension expenditure for 2005 to 2007, and for the latter, I calculate the average contribution from 2012 to 2014. Then, for all 198 cities for which I have at least one year of data in the early period and

one year of data in the later period,¹¹ I calculate the percentage change between the two. For one of the 198 cities—Fresno, California—the percentage change is undefined because the city contributed \$0 to its pensions in each year from 2005 to 2007. I also exclude Birmingham, Michigan: due to contributions at or near zero for the first three years, its pension contributions have increased 2,151% over the ten-year period. For the remaining 196 cities, however, this statistic nicely captures how within-city pension costs have changed over time.

Figure 2 shows the distribution of this variable. Strikingly, the number is positive for nearly all of the cities: 165 cities have seen increases in their pension expenditures from 2005-2007 to 2012-2014, and only 31 have seen decreases. What is more, among the cities with decreases are cities like Chicago and Dallas, as well as San Diego, which made a very large one-time payment in 2006 to settle several pension-related lawsuits. While there is considerable variation in the percentage change in pension expenditures across these cities, with a standard deviation of 63%, the magnitude of the increases is quite large in many cities. The median increase is 30%, and the average is 47%. More than a quarter of the cities have weathered increases larger than 68%. And in 25 of the cities—about 13% of the sample—pension costs have more than doubled within cities over this time period. Thus, it does not appear that pension costs are only growing in a small set of cities. Instead, rising pension costs appear to be the norm in U.S. municipal governments.

Of course, if pension expenditures have been steadily increasing, then by averaging annual pension contributions from 2005 to 2007 and 2012 to 2014, I am actually understating the extent of the increases. In Figure 3, therefore, I look strictly at the percentage change in

¹¹ There are 180 cities for which I have data for all of these years, 2005-2007 and 2012-2014. For an additional 18 cities, I am missing either one or two years from the early period or one or two years from the later period. I include these cities as well using the years of data that I have.

contributions from 2005 to 2014 for the 181 cities for which I have contributions data in both of those years. As before, I find that 85% of the cities have experienced increases. But here I find that the median within-city change is 45%, and the average is 69%. A total of 45 cities—almost a quarter of the sample—have seen their pension expenditures more than double in ten years. Clearly, then, almost all cities are experiencing growth in their pension expenditures.

The analysis so far has focused on changes in pension costs within cities over time, but the *level* of pension expenditures also varies dramatically across cities—and not just because of differences in city size. Returning to Figure 1, note that Coralville, Iowa, and Eastlake, Ohio, are roughly the same size, but median annual pension expenditures in Coralville are \$605,860, while they are \$1,015,945 for Eastlake. Buffalo, New York, is actually smaller in size than Anchorage, Alaska, but median pension expenditures in Anchorage are \$31,866,160, while they are \$53,183,700 in Buffalo. In order to make pension expenditure figures comparable for all cities in the dataset, I calculate annual city pension expenditures per full-time-equivalent city employee using data from the Census’s Survey of Governments Employment statistics. Of the 2,081 city-year observations of pension expenditures in my dataset, which come from 219 unique cities, the Employment dataset includes information on 1,748 city-years in 216 unique cities.¹²

For purposes of illustration, I take the within-city median of this variable, sort those medians from smallest to largest, and display in Table 1 the value of the medians for every fifth city in the dataset. Clearly, there is tremendous variation in per-employee pension expenditures across these 216 cities. Cities in the middle range spend about \$7,000 per employee per year, but municipalities such as Norman, Oklahoma, and Parsons, West Virginia, spend less than \$2,500,

¹² Most of the city-years that do not find matches in the Census Employment statistics are from the earliest years, 2005 and 2006.

while Westland, Michigan, Warren, Michigan, and San Diego, California, spend more than \$20,000. Within the full dataset of 1,748 cases, the median per-employee annual pension expenditure is \$7,175, but the standard deviation is very large: \$16,234 per employee per year. In fact, *most* of the variation in per-employee pension expenditures in this dataset is across cities: an OLS regression of logged annual per-employee pension expenditures on a set of city fixed effects has an R-squared of 0.85 (not shown).

What explains variation in pension costs across cities? Based on the discussion above, one possibility worth exploring is whether pension expenditures per employee are higher in cities with collective bargaining and strong public-sector unions. If pension benefit levels are higher in unionized cities, as some scholars have argued, that would tend to increase cities' required pension contributions in strong-union states relative to weak-union states. Also, if pension underfunding is more severe in states with strong unions, that might affect cities' pension expenditures—although it is not clear in which direction. If cities in heavily unionized states have to contribute more toward pensions to make up for greater funding shortfalls, their contributions should be larger than in less unionized states. But if city contributions themselves reflect a greater tendency to underfund, then city expenditures might actually be lower in heavily unionized states. It is an empirical question, then, what overall effect collective bargaining and strong unions should have on a city's pension expenditures.

Beyond a city's collective bargaining status and unionization rates, the patterns of Table 1 suggest that larger or more urban cities might have greater pension expenditures per employee, perhaps because of higher cost of living. In addition, DC plans are often promoted as less expensive for governments than defined benefit plans, and so if cities are more reliant on DC plans than defined benefit plans, we might expect their overall retirement expenditures to be

lower. In addition, Munnell (2012) proposes that cities with larger numbers of plans, or with more locally-operated (as opposed to state-operated) plans, may have more generous pension benefits, which would tend to push up pension expenditures per employee. And of course if a city's pension expenditure figures include OPEB or payments from POB revenue, that would also increase the contribution amounts for that city.

To explore these matters empirically, I regress logged per-employee pension expenditures for the full dataset on a set of state-, city-, and pension-related variables. The first is an indicator for whether the city is in a state that requires government employers to engage in collective bargaining 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 as a 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. Second, I include logged city population and percent of the city that is urban from the 2010 U.S. Census. I also include logged income per capita from the American Community Survey (the five-year estimates from the 2013 survey, adjusted to 2014 dollars). The other predictors come from the pension dataset: the city's total number of plans, percent of the city's total retirement expenditures that went to DC plans, percent of contributions that went to locally-operated plans, an indicator for whether the city's retirement expenditure figures include OPEB, and an indicator for whether the city issued POBs that year. I cluster the standard errors by city to correct for autocorrelation within cities over time.

Column 1 of Table 2 presents the estimates of this model.¹³ As expected, there are systematic differences in per-employee pension expenditures in cities of different sizes. On

¹³ I lose one observation by logging the dependent variable (\$0 in contributions). I also drop two observations from Laramie, Wyoming, because it spent over \$400,000 per employee on pensions in 2007 and 2012, and those two observations have a large effect on the coefficient estimates.

average, a 10% increase in city population is associated with a 0.8% increase in annual pension expenditures per employee. The coefficient on income per capita is not significant, but the coefficient on percent urban is: a shift from a rural to an urban area is associated with a 37% increase in pension expenditures. However, the results give no indication that having a larger number of plans or more locally-operated plans have significant effects on pension expenditures. The coefficient on the share of DC contributions is also insignificant, suggesting that cities more reliant on DC plans are not necessarily spending less on retirement overall.¹⁴ The coefficient on the POB indicator, however, is positive and significant, as is the coefficient on OPEB inclusion.

But the most striking finding in column 1 of Table 2 is the coefficient on the collective bargaining indicator: it is positive and statistically significant at the 1% level. On average, cities in states with duty-to-bargain laws for local employees spend 40% more on pensions per employee per year. Thus, even controlling for differences in city size, income, and urban population, collective bargaining is associated with significantly higher pension expenditures.

Because duty-to-bargain laws are so strongly correlated with public-sector union membership in a state (see Saltzman 1985, 1988, Moe 2011)—a correlation of 0.84 in this dataset—there is little to be gained from estimating separate effects of unionization and collective bargaining.¹⁵ In column 2, however, I replace the collective bargaining indicator with the percentage of state and local government employees in the state who are in unions.¹⁶ Here,

¹⁴ However, only 58 of the cities made contributions to DC plans, and for those 58, the median share of total contributions going to DC plans is only 33%.

¹⁵ When I include both the duty-to-bargain indicator and the unionization measure in the model, both have positive coefficients, but only the duty-to-bargain indicator is significant.

¹⁶ These figures are Anzia and Moe's calculations using Current Population Survey data from 2000 to 2010. They include only full-time government employees who work at the level of state or local government. It is constant within states for all years of the dataset.

too, I estimate a positive, significant effect of unionization. On average, a 10% increase in union membership is associated with a 7% increase in city pension contributions per employee.

Finally, in column 3, I return to the duty-to-bargain indicator and account for average over-time growth in pension expenditures by including fixed effects for each year. Mirroring the within-city patterns described earlier, the coefficients on the year indicators (not presented) are steadily rising, making it clear that pension expenditures have risen over time. But even accounting for this over-time growth, cities in duty-to-bargain law states spend much more per employee per year on pensions: 40% more on average.

By simply collecting and presenting data on public pension expenditures in a large set of cities, therefore, we learn a great deal that is new. Scholars and practitioners have made conflicting assertions about what has been happening to local governments' pension costs over time, and about how heavy a burden pension costs are for the typical local government. Some have argued that local governments' pensions are tremendously expensive and that their costs have increased dramatically in recent years—reaching crisis levels across the United States. Others have maintained that cumbersome, rapidly-rising pension costs are only a problem for a few local governments. With these new data in hand, we can actually evaluate these competing claims. Based on what I have found so far, it is true that local pension expenditures have risen—just about everywhere. But there is also significant variation across cities in both their levels of pension expenditures and the degree to which their pension expenditures have increased. Both kinds of variation stand to affect what cities have done—and have had to do—to adapt.

Are Pensions Crowding Out Public Services?

I turn next to the question of whether rising pension costs are crowding out public services. As a practical matter, this requires a definition of public service provision that can be

operationalized for all or most city-years in the pension costs dataset. While there are many possible ways of thinking about public service provision—including dimensions such as frequency, geographic scope, equity, or quality—many of these dimensions are not easily measured for a large number of cities over multiple years. One feature of city government that *can* be measured, however, is city employment.

City employment levels are also the most natural place to begin looking for pension-related service reductions, for reasons beyond data availability. More so than at the national level, the amount and quality of local government service provision is heavily dependent on the employees providing the services: the police providing law enforcement, the firefighters providing fire protection, the sanitation workers cleaning streets and collecting refuse, and so on. Also, a city's pension costs are in part a function of their employment levels: as a city hires more employees, it has to begin making pension contributions for them. Therefore, city officials under pressure to reduce spending to make way for rising pension expenditures might find reducing city employment an attractive (or necessary) option. In addition, anecdotal evidence from a few high-profile cases demonstrates that some cities have indeed responded to rising pension costs by trimming the ranks of their public workforces. Therefore, city employment levels are a natural place to begin looking for effects of rising pension costs.

The question, then, is straightforward: Do city officials respond to rising pension costs by reducing city employment? Setting up a model to test the effects of rising pension costs on city employment, however, is less straightforward than it might seem, the reason being that—as I've just explained—city pension contributions are in part a function of the number of city employees. That means that the rising pension expenditures documented in Figures 1-3 do not

necessarily portend trouble for municipal service provision. They could, in fact, be a sign of *improvement* in government service provision if cities are expanding their workforces.

Figure 4 provides a helpful illustration. There, for nine cities in the dataset that I've selected, I plot the yearly values of two variables: total pension expenditures in 2014 dollars (labeled on the left-hand vertical axis) and total full-time-equivalent employment (labeled on the right-hand vertical axis). In each of these cities, pension expenditures have been rising. But so has city employment. For cities of this type, then, a regression of city employment on pension expenditures would produce a positive coefficient on pension expenditures—but not because rising pension costs are leading city officials to increase employment. Rather, the positive coefficient would reflect that pension expenditures *should* go up when cities hire more employees, because the cities are making pension contributions for a larger number of people. The default expectation, then, is a positive relationship between these variables.

That makes the pattern in Figure 5 all the more striking. There, I plot the same two variables over time for another set of nine cities. Just like the cities in Figure 4, pension expenditures in these cities have steadily increased over time, but employment levels have *decreased*. For example, in Sunnyvale, California—a city of 140,000 people—pension expenditures have increased from \$13.9 million in 2005 to \$24.2 million in 2014, in real terms. Yet full-time equivalent employment has dropped over 20% in the same time period, from 1,100 employees in 2005 to only 872 in 2014. Given the default expectation that these variables should be positively related, the negative relationship in Figure 5 is especially remarkable. It suggests that there very likely *is* crowd-out in these cities. It also suggests that what is happening in cities like Sunnyvale is probably very different than what is happening in places like Bozeman, Montana, and Lawrence, Kansas—cities depicted in Figure 4.

This discussion also highlights the complexity of estimating the effects of pension expenditures on employment levels when pension expenditures themselves are partially a function of employment levels. I propose two possible modeling approaches for dealing with that complexity. Both involve lagging the independent variables by one year so that, in any model, I am estimating the effects of pension costs in year $t-1$ on employment levels in year t . This helps to address the mechanical endogeneity of pension costs and employment in the same year, but it also models government decision-making in a realistic way: officials presumably make decisions about next year's expenses based on what they observe of this year's. However, the two approaches differ in the operationalization of the main independent variable and in the theoretical expectations for the coefficient estimates. The first approach involves modeling city employment as a function of total city pension expenditures. In this approach, if the estimated coefficient on pension expenditures is anything other than positive, one could conclude that in the average city, rising pension costs are *not* simply a reflection of workforce expansion. The second approach involves modeling city employment as a function of pension expenditures *per employee*. This approach allows a more direct assessment of service crowd-out: An insignificant coefficient would imply that in the average city, rising per-employee pension expenditures are *not* associated with employment reductions. A negative coefficient, however, would imply that greater increases in per-employee pension costs *are* causing employment reductions—evidence that services are being crowded out.

Both approaches are reasonable, and both allow me to test whether rising pension expenditures in the average city are a benign response to workforce expansion or a trend that is occurring even as employment levels are stagnant or decreasing. That said, the second approach lends itself toward a more direct interpretation of the effects of pension costs on service crowd-

out, and so that is the approach I adopt in the analysis to follow. The main independent variable, then, is logged per-employee pension expenditures in each city-year.

The dependent variables come from the U.S. Census's Survey of Governments Employment files, which have information on full- and part-time city employment and payrolls for 1,748 of the city-years in my pension costs dataset. Because large cities naturally have more employees than small cities, I divide each of the employment figures by city population as of the 2010 Census. The distributions of these per-capita city employment variables have long right tails, so I take the log of each one.

A first pass at the data shows that most of the variation in per capita employment is across rather than within cities. Thus, even if there is a negative bivariate relationship between city pension costs and employment levels, that may not tell us about how cities are *changing* their employment levels in response to rising pension costs. In the analysis to follow, therefore, I use OLS with city fixed effects. City fixed effects partial out the influence of any time-constant characteristics of cities that lead them to have higher or lower employment levels and pension expenditures—and thus put the emphasis on changes within cities over time.

It is less clear whether year fixed effects are warranted. My primary interest is in the effects of rising pension expenditures, which—as I have shown—have risen steadily in almost all cities over the ten-year period from 2005 to 2014. By including year fixed effects, I would partial out most of the effects of the very trend I am interested in. However, there are secular trends that likely affect pension costs and employment in all cities, such as the Great Recession. Moreover, some cities have had greater pension expenditure increases than others, and including year fixed effects allows me to test whether those greater-than-average increases are associated with greater-than-average reductions in employment. In what follows, therefore, I present the

results of models both with and without year fixed effects. I also include logged city general revenue per capita, with the expectation that revenue will be positively associated with employment. This is an important control variable because any negative relationship between pension expenditures and employment could partially be an effect of the Great Recession: during those years, required pension expenditures increased because of the decline in asset values, and at the same time, city revenues dropped—perhaps leading to lower employment.

Most of the variables from my earlier analysis (in Table 2) are either constant within cities over time or do not have any relationship with city pension costs per employee, so I do not include them in these models. The exception is the indicator for whether the city issued POBs. New POBs clearly are associated with higher pension contributions, but they may also be positively correlated with employment, in that most cities likely issue POBs so that they do not have to reduce employment or cut expenditures. Thus, in the models, I include a dummy variable equal to one for each city-year in which new POBs were issued.¹⁷

In column 1 of Table 3, I start by presenting the results from the most basic model: the dependent variable is logged full-time equivalent employment per capita, and the independent variables (all lagged by one year) are logged per-employee pension expenditures, logged general revenue per capita, the POB indicator, and the city fixed effects. The coefficient estimates in column 1 suggest that rising pension expenditures have indeed led to reductions in city employment in the average city: a 10% increase in city pension expenditures per employee is

¹⁷ Because I am now analyzing changes within cities over time, I exclude cities from the analysis if they have fewer than three years of pension cost and employment data. I also exclude one observation from the analysis because it is a leverage point: Detroit in 2014, because its pension expenditures dropped from \$6,983 per employee in 2012 to \$750 per employee in 2013. However, the main reason the number of observations here is lower than in Table 2 is that I am lagging the independent variables by one year.

associated with a 0.73% decrease in overall city employment the following year—an effect that is significant at the 1% level. In addition, the effects of general revenue and the POB indicator are positive, as expected. This, then, is early evidence that rising pension expenditures per employee are leading to service reductions in the typical city.

When I add year fixed effects to the model, as I do in column 2, the effect of rising pension costs remains negative and statistically significant. The coefficient is smaller in magnitude than in column 1, which is to be expected given that the year dummies partial out the effects of average over-time increases in pension expenditures. Even so, I find that a 10% increase in pension costs per employee is associated with a significant 0.43% decrease in employment per capita the following year. Thus, cities with larger-than-average pension cost increases year-to-year have experienced larger-than-average reductions in city employment.

If cities are reducing employment in response to rising pension costs, where should we expect to see the greatest decreases? One hypothesis is that the cuts should be greater among full-time city employees than part-time city employees, because part-time employees often aren't eligible for pensions. I explore this in columns 3-6 of Table 3. In columns 3 and 4, the dependent variable is the log of the city's *full-time* employees per capita, which again I model without year fixed effects (column 3) and then with year fixed effects (column 4). In both columns, the coefficients on log per-employee pension expenditures are negative and statistically significant. When I instead model the log of the city's *part-time* employees per capita, in columns 5 and 6, the coefficients on pension costs are negative but statistically insignificant.¹⁸

¹⁸ The number of observations is lower here due to a few cities without part-time employees.

What this suggests is that on average, growing pension costs are associated with declining employment levels—but for full-time employees, not necessarily part-time employees.¹⁹

This is one way to begin exploring heterogeneity in the impact of rising pension costs, but even if we focus solely on full-time employees, cities might respond very differently to the same percentage increases in their pension expenditures. One potentially important difference is whether the city is in a state with collective bargaining and strong unions. Table 2 showed that cities' pension expenditures per employee are 40% higher in states that have duty-to-bargain laws. Because of this, the same percentage increase in city pension expenditures might have a bigger effect on city budgets in cities with collective bargaining—requiring bigger employment reductions. In addition, cities' collective bargaining contracts and politically active unions might limit officials' options for cost reduction. For example, compared to cities with weaker unions, officials in cities with strong unions and collective bargaining might have a harder time opposing salary increases, larger employee contributions, or reducing other fringe benefits. If so, then perhaps we should expect pension-induced employment reductions to be more pronounced in cities with collective bargaining and strong unions.

I test this hypothesis in column 1 of Table 4. There, I interact pension costs with the indicator for state collective bargaining laws—thereby allowing the effect of increasing pension costs to vary with the state's bargaining provisions. The coefficient on the pension expenditures variable remains negative, but it is smaller than before and statistically insignificant at the 10% level. This means that for the average city in a state *without* a duty-to-bargain law, larger pension expenditure increases are not significantly associated with larger decreases in full-time

¹⁹ Cities could also be reducing the hours worked by part-time employees, even if the number of part-time employees has stayed the same. However, when I model logged part-time employee *payroll* in each city, which reflects total hours, I find no significant effect of pension costs.

employment. For cities in duty-to-bargain law states, however, the effect is large and negative: the combined coefficient for log per-employee pension expenditures and its interaction with the collective bargaining law indicator is -0.099 (see the bottom of the table). This means that for the average city in a duty-to-bargain law state, a 10% increase in per-employee pension costs is associated with about a 1% decrease in full-time city employment. When I add year fixed effects in column 2, my conclusion is largely the same. There is no clear effect of larger-than-average increases in pension costs in non-bargaining states, but there is in bargaining states: specifically, a 0.66% decrease in full-time employment for every 10% increase in pension expenditures.²⁰

Could these models simply be picking up differences in the effects of pension cost increases in urban and rural municipalities? After all, more of the cities in duty-to-bargain law states are urban, and urban municipalities also have higher per-employee pension costs than rural municipalities (see Table 2).²¹ There aren't many rural municipalities in this dataset—17 in duty-to-bargain states and 24 in non-duty-to-bargain states—but even so, in columns 3 and 4 of Table 4, I add an interaction of city pension costs and an indicator for urban cities, and I also add a triple interaction of pension costs, duty-to-bargain status, and the urban indicator. The results suggest that pension-induced employment reductions have indeed been larger in urban than in rural municipalities—but also that the presence of collective bargaining makes a meaningful difference. In both columns 3 and 4 (the models without and with year fixed effects, respectively), I find no significant effect of rising pension costs on employment levels for rural municipalities—regardless of whether they are in duty-to-bargain law states. Column 3 does

²⁰ The effects are substantively the same when I replace the duty-to-bargain law indicator with state public-sector union membership: the negative effect of pension costs on full-time employment grows in magnitude as the unionization rate increases.

²¹ Of course, urban municipalities in non-duty-to-bargain states are also more likely than rural municipalities to have active public-sector unions.

suggest that the effect of rising pension costs has been negative for the average urban city in a non-bargaining state, but that effect is insignificant when I add year fixed effects to the model (column 4). For urban cities in bargaining states, though, the effect of rising pension costs is clearly negative for full-time city employment, regardless of the specification. In column 3, I estimate that a 10% increase in pension costs leads to a 0.94% reduction in full-time employment, and in column 4, the effect is a 0.56% reduction.

As a final step in my analysis of city employment, I explore whether certain *types* of employees are more affected than others by rising pension costs. While each city has its own mix of employees, and cities may well reduce employment in different ways, most cities do have public safety employees—police officers and firefighters—and in the typical city, public safety workers make up a large share of all full-time city employees: 37% on average in this dataset. Also, police and fire protection employees often have their own pension funds, sometimes with more generous benefits than non-safety workers. Both factors might make public safety employees more vulnerable to pension-related employment cutbacks than non-safety workers. At the same time, public safety employees provide core local government services, and they also tend to be well-organized and very active in city politics (Anzia and Moe 2015). Thus, even though they are a large part of the typical city workforce and often have their own pensions, it is possible that they are *more* protected from cutbacks than non-safety workers.

To test this, in Table 5, I examine three categories of city employees: fire protection employees, police protection employees, and non-safety and non-education employees. (Only 21 of these cities handle education, but in those cities, education employees are a large share of the total city workforce, so I exclude them.) I continue with the model from column 1 of Table 4, where I allow the effect of rising pension costs to vary depending on whether the city is in a

duty-to-bargain law state. Therefore, I am testing whether increases in pension costs are associated with decreases in full-time employment of fire protection employees, police protection employees, and non-safety employees—and whether there are different effects for each for cities in states with and without collective bargaining laws.

I begin with the public safety employees. Interestingly, in columns 1 and 2, I see no discernable effect of pension cost increases on fire protection employment, regardless of whether the city is in a state with a collective bargaining law.²² Thus, it does not appear that rising pension expenditures are leading to decreases in the number of firefighters in the typical city. For police protection employees, however, the results are different. In column 3, the model without year fixed effects, I find that pension contributions are associated with reductions in police protection employees in collective bargaining states: specifically, a 10% increase in per-employee pension expenditures is associated with a 0.58% reduction in full-time police protection employment. This is consistent with the results of Anzia and Moe (2015), who find evidence that police officers in American cities confront a steeper tradeoff between employment levels and per-employee compensation than firefighters do. That said, when I add year fixed effects in column 4, the effect for police employment shrinks in magnitude and is no longer significant at conventional levels ($p=0.159$). Taking these results together, then, it appears that public safety workers are relatively protected from pension-induced cuts in the average city, although there may be modest negative effects for police protection employees.

In columns 5 and 6, I examine non-safety (and non-education) employees, and the patterns are much clearer—and much more negative. The coefficient estimates show that

²² There are fewer observations in this model than in earlier models because not every city has fire protection employees.

overwhelmingly, this is where the pension-induced employment cuts are happening in the average city. In column 5, I find that even in cities without collective bargaining, a 10% increase in pension expenditures per employee is associated with an average 0.46% decrease in non-safety employment per capita. But the larger effects, again, are in cities in collective bargaining law states. In those cities, a 10% increase in pension expenditures is expected to reduce non-safety employment by 1.56%, and that effect is significant at the 1% level. In model 6, where I add year fixed effects, the effect on non-safety employment is statistically insignificant for cities in non-duty-to-bargain states, but the effect for cities in collective bargaining states remains large and negative. Thus, if public safety employees have been relatively buffered from employment cuts related to rising pension costs, non-safety employees have not.

The employment reductions that have been occurring in many cities stand as evidence that public services are indeed being reduced alongside rising pension costs. In a final round of analysis, I move beyond employment as the indicator of service crowd-out and test for the effects of rising pension costs on capital outlays.²³ This is another important dimension of city activity because it relates to the city's investments in construction and the purchase of land, equipment, and existing buildings. 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), although capital outlays are by no means limited to those functions.

²³ The difficulty in using many of the city spending variables for this analysis is that the Census's Survey of Governments Finance files do not make it clear which variables include pension costs, and it would be problematic to try to examine the effects of rising pension costs on expenditure variables that might include those very pension costs. One of the nice features of both the employment (and payroll) figures as well as the capital outlays figures is that I can be confident that they do not include pension contributions.

Therefore, this final analysis assesses the effects of rising pension costs on a spending area that is separate from employment and payroll expenditures.

In Table 6, I regress logged per capita city outlays on the same variables as in the earlier tables. In column 1, the model that includes city fixed effects but not year fixed effects, I find a pattern that mirrors the patterns I found in the employment analysis: For cities in states without collective bargaining laws, as per-employee pension costs have increased, spending on capital projects has decreased—but the effect is not significant at conventional levels. But for cities in states with bargaining laws, the effect is large, negative, and statistically significant. The estimates show that within cities, a 10% increase in per-employee pension costs is associated with a 3% decrease in capital outlays per capita. This pattern also holds in the model with year fixed effects. In column 2, the effect of a 10% increase in pension costs is a 2.1% reduction in capital outlays, significant at the 10% level. 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—yet another way in which pension costs in some places are crowding out services.

Conclusion

Nearly all scholarly work on public 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 argue that generous pension benefits combined with severe underfunding are putting pressure on the budgets of local

governments—forcing them to cut back on public services. Others maintain that the big problems are limited to a few governments that have been fiscally irresponsible. Until now, there haven't been any data available to assess which is closer to the mark.

This paper introduces a new dataset that allows us to evaluate where pension expenditures tend to be highest, the extent to which they have increased over time, and whether they are crowding out public services. I collected up to ten years of annual financial documents for a diverse set of 219 cities throughout the country, and then I combed each one for information about what they spent on their employees' pensions. A simple descriptive analysis of these data reveals a great deal that is new. City pension costs *are* rising—almost everywhere. For the average city in the dataset, inflation-adjusted pension expenditures rose 69% between 2005 and 2014, and in about a quarter of the cities, pension expenditures have more than doubled. There is also wide variation in the level of pension expenditures across cities: large, urban areas and cities in states with collective bargaining laws spend much more per employee on pensions than smaller, rural areas and cities in states without duty-to-bargain laws. Moreover, I find that rising per-employee pension expenditures are associated with declines in service provision—but only for certain cities and certain employees. Looking at both employment and capital outlays, I find that the impact of rising pension expenditures has been most negative for cities in states with collective bargaining laws. But certain kinds of employees, such as public safety employees, have been relatively protected in the average city: employment reductions appear to be most pronounced for full-time non-safety employees.

As I said earlier, these data cannot be used to evaluate which cities are “in trouble,” because the pension expenditure figures tell us what cities are *actually* spending on pensions—not what they *should* be spending or how underfunded their pensions are. This is an important

caveat. It remains possible that certain cities with low per-employee pension expenditures today will experience big problems down the road if their expenditures are so low *because* they are underfunding. At the same time, these data are much more useful than plan- or state-level data for assessing what cities are experiencing and how they have been adapting to pension increases.

Going forward, it will be important to build on this data collection by extending it to a larger number of municipal governments as well as counties, school districts, and special districts. The data collection approach is costly and time-consuming—which is why I only selected about 200 cities—but there is little reason to think that the pension cost increases shown here are limited to municipal governments. To the contrary, nearly all full-time local government employees are eligible for defined benefit pensions, and anecdotal evidence suggests that pension expenditures are also an increasing burden for counties and school districts.²⁴ Future research should study this closely—and should carry out evaluations of whether pension costs are affecting service provision in these kinds of governments as well.

In closing, I return to the question I began with: Have state and local governments in the U.S. entered a New Fiscal Ice Age—an era in which local governments provide less service at higher cost (Kiewiet and McCubbins 2014)? Or are major pension problems limited to only a few cities that have been especially irresponsible? To some extent, the answer is yes to both. Pension expenditures *are* rising almost everywhere. In some places, however, those increases are partially a natural consequence of rising employment—and therefore may even indicate service improvements. In the average city, however, pension cost increases have come hand-in-hand with *reductions* in city workforces and *decreases* in city capital expenditures. In the latter

²⁴ See, e.g., Jill Tucker, “California schools may face cuts amid skyrocketing pension costs,” *San Francisco Chronicle*, January 23, 2017, <http://www.sfchronicle.com/bayarea/article/California-schools-may-face-cuts-amid-10873046.php> (accessed July 31, 2017).

cities, it looks as though rising pension expenditures *are* changing the landscape. And to the extent that this transformation is taking place, it is no small change. It stands to affect employment opportunities in the public sector—and roughly 15 million people, or about 6% of the adult population, work for U.S. state and local government on a full-time basis. But more generally, it suggests that citizens of many cities can anticipate a future in which their local governments do less with more. This analysis only provides a first look at what is happening to city pension costs across the United States, but the picture that emerges is not promising—and it suggests that the future of local government may look very different than the past.

References

- Andonov, Aleksandar, Yael V. Hochberg, and Joshua D. Rauh. 2016. "Political Representation and Governance: Evidence from the Investment Decisions of Public Pension Funds." Working paper.
- 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. 2017a. "Polarization and Policy: The Politics of Public-Sector Pensions." *Legislative Studies Quarterly* 42 (1): 33-62.
- Anzia, Sarah F., and Terry M. Moe. 2017b. "Interest Groups on the Inside: The Governance of Public Pension Funds." Working paper.
- DiSalvo, Daniel. 2015. *Government Against Itself: Public Union Power and Its Consequences*. New York: Oxford University Press.
- 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.
- 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.
- 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.
- Pew Charitable Trusts. 2016. "The State Pension Funding Gap: 2014." Issue brief, August 24, 2016.
- Rauh, Joshua. 2016. "Hidden Debt, Hidden Deficits: How Pension Promises are Consuming State and Local Budgets." *Hoover Institution Essay*.

- Romano, Roberta. 1995. "The Politics of Public Pension Funds." *Public Interest* 119: 42-53.
- 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.
- 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.

Figure 1: Total pension contributions by city, 2005-2014

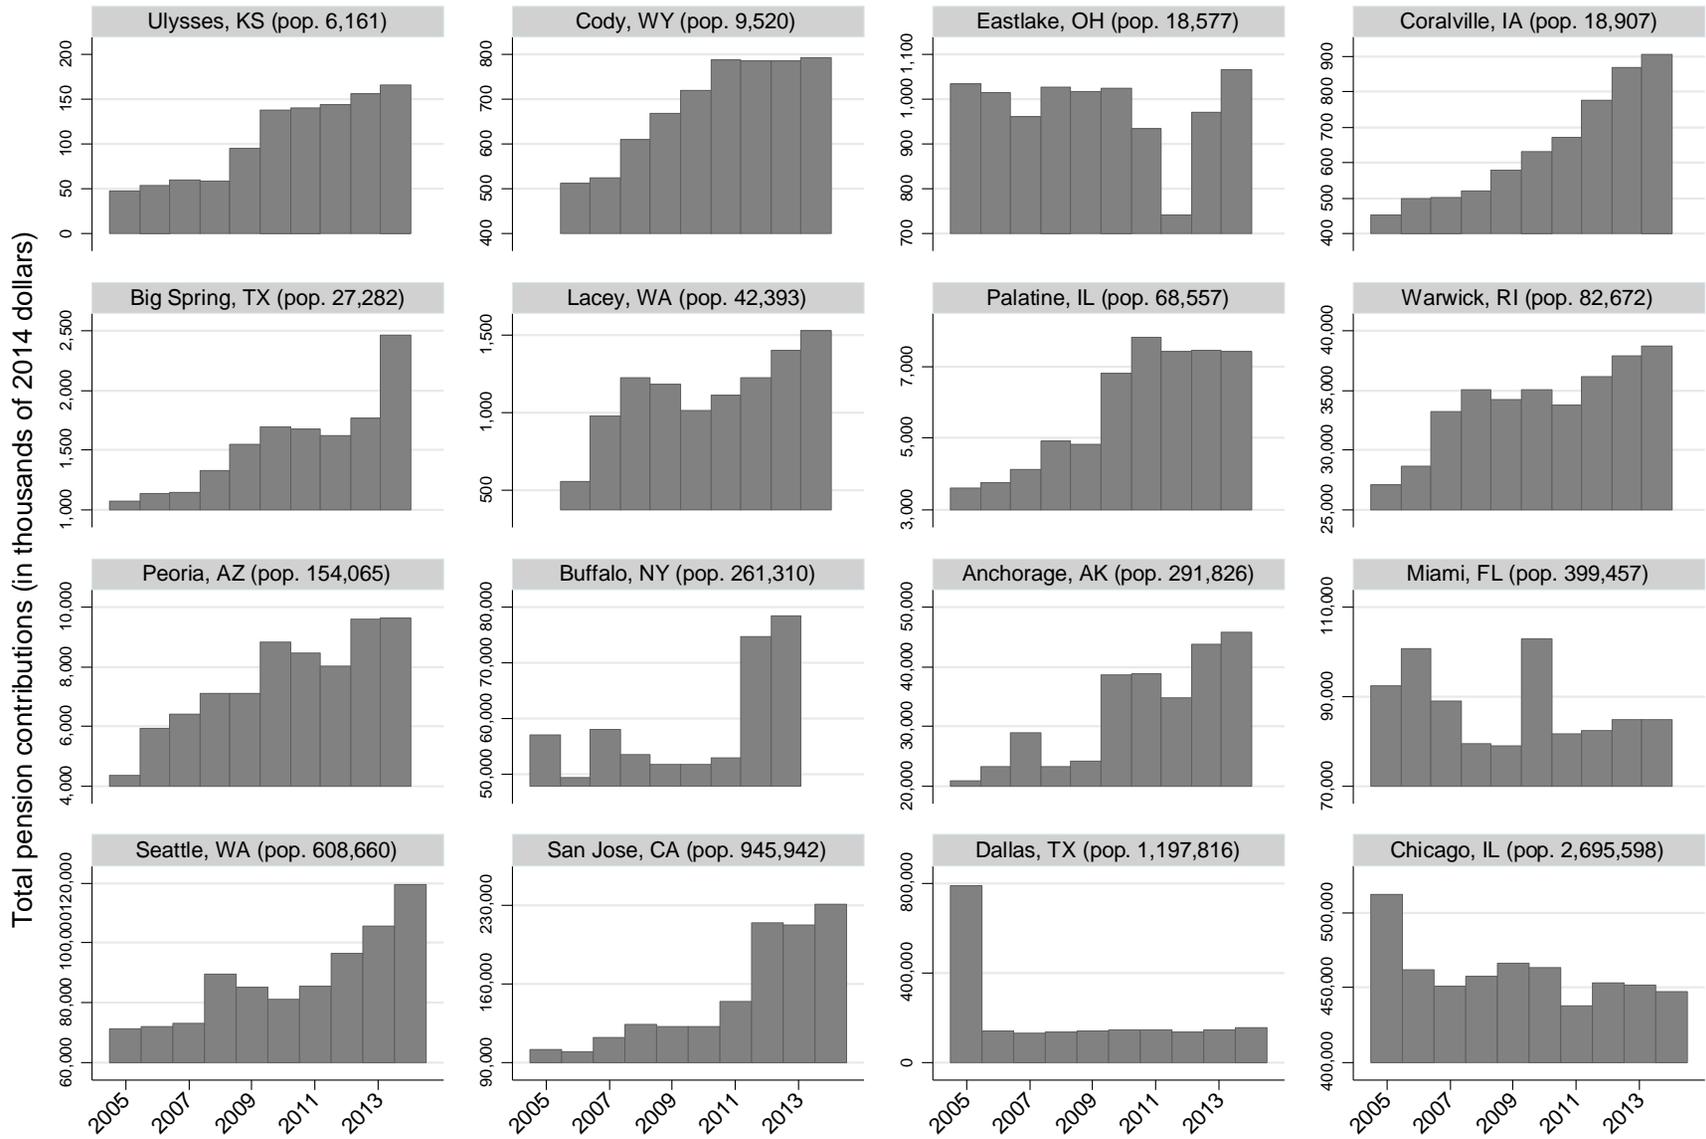


Figure 2: Percent Change in Annual Pension Expenditures, from 2005-2007 to 2012-2014

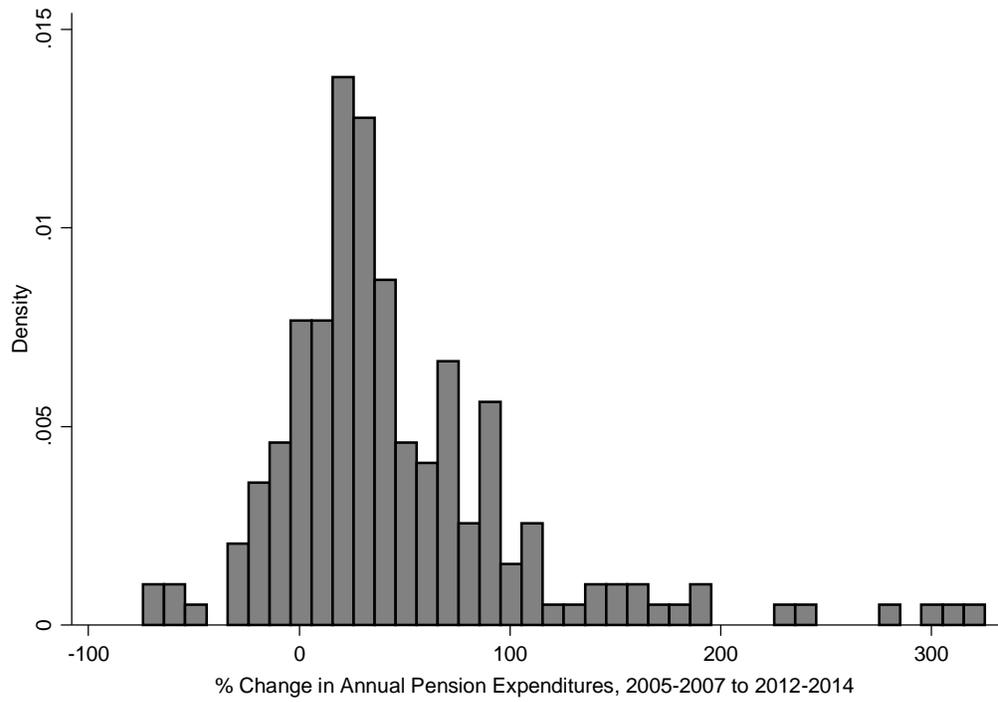


Figure 3: Percent Change in Annual Pension Expenditures, from 2005 to 2014

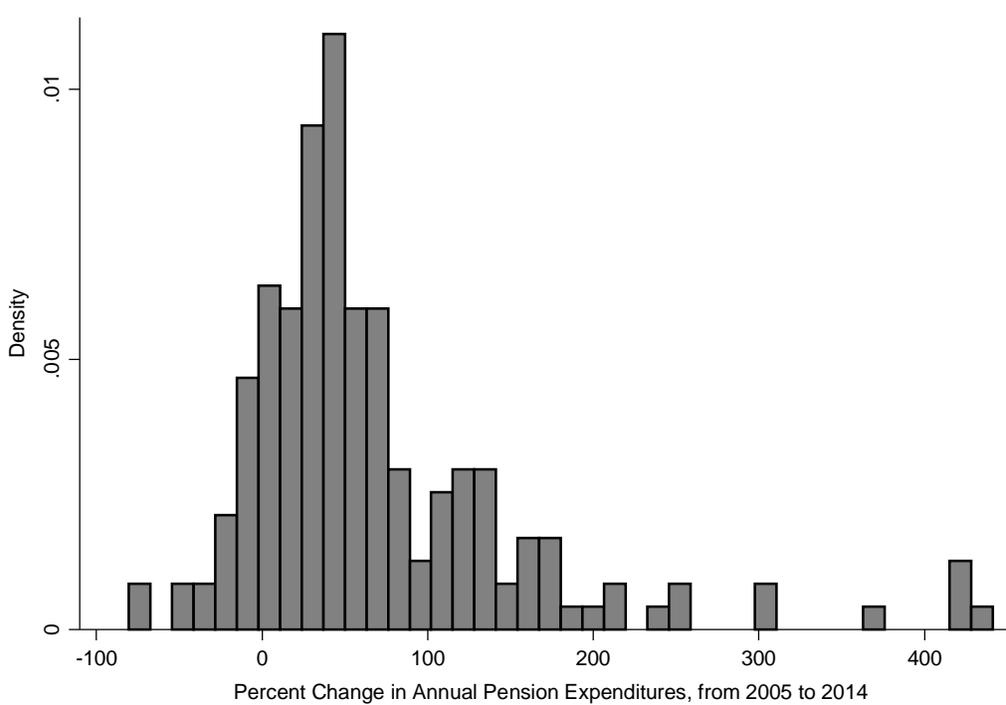


Figure 4: Pension expenditures and total employment, positive relationship

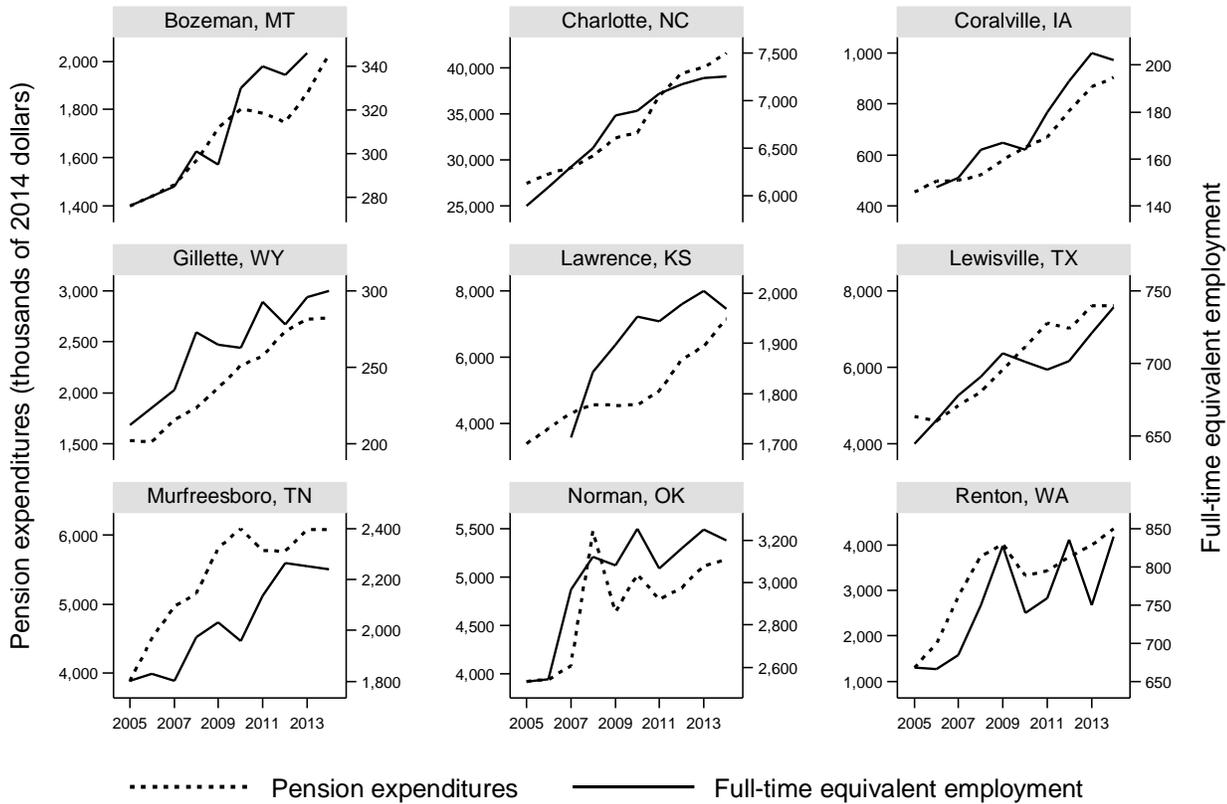


Figure 5: Pension expenditures and total employment, negative relationship

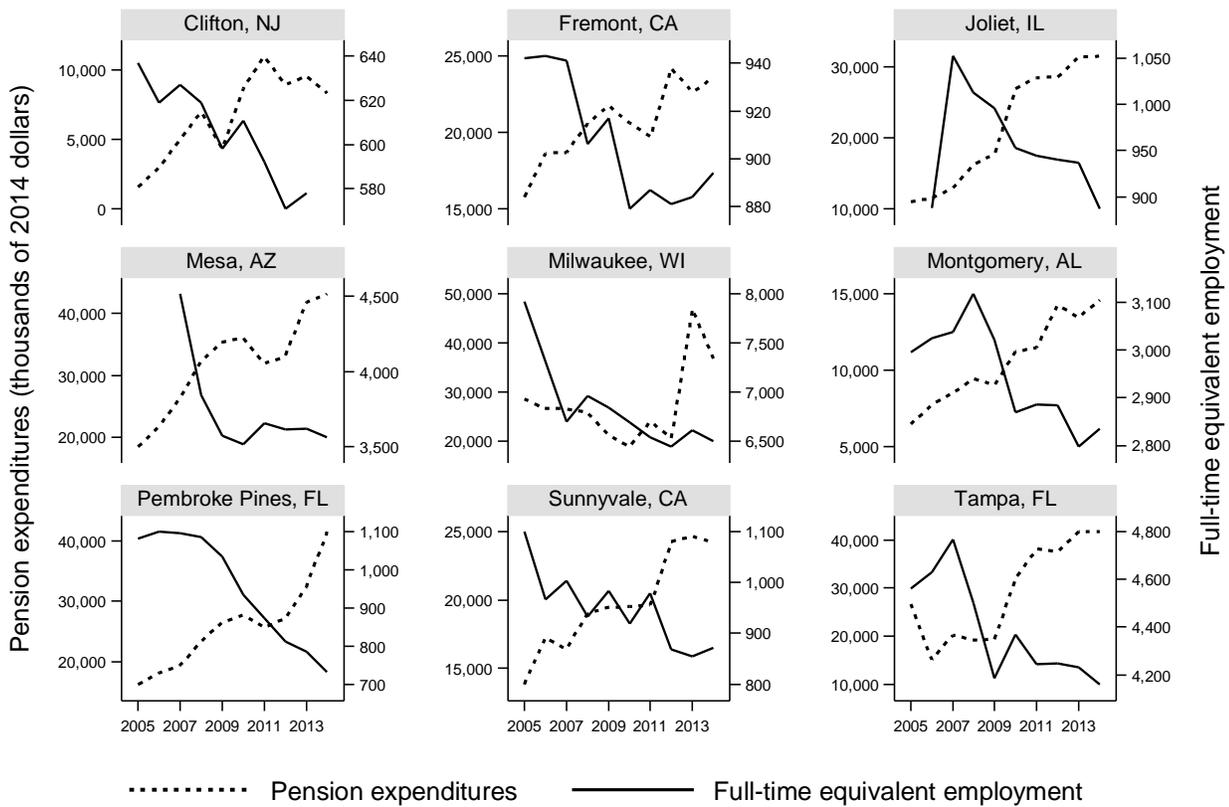


Table 1: Median annual pension expenditures per employee (2014 dollars)

Norman, OK	1,549	Dickson, TN	7,551
Parsons, WV	2,213	Corpus Christi, TX	7,762
Murfreesboro, TN	2,703	Gillette, WY	8,060
Brookings, SD	2,961	Eagan, MN	8,153
Williston, ND	3,356	Biloxi, MS	8,300
High Point, NC	3,698	Noblesville, IN	8,709
Essex Junction, VT	3,809	Honolulu, HI	9,125
Riverton, WY	4,027	Minneapolis, MN	9,355
Columbia, SC	4,173	Mesa, AZ	9,547
Renton, WA	4,512	St. Louis, MO	9,837
Wilson, NC	5,076	Tupelo, MS	10,032
Kent, WA	5,222	Woodstock, IL	10,396
New Bern, NC	5,448	Ft. Worth, TX	10,794
Norfolk, VA	5,712	Midland, TX	11,231
Baltimore, MD	5,896	Phoenix, AZ	12,306
Fresno, CA	6,066	Newark, NJ	13,661
Highland Park, NJ	6,220	Chula Vista, CA	14,675
Fort Dodge, IA	6,302	San Jose, CA	18,089
Cody, WY	6,481	Westland, MI	22,191
Morristown, TN	6,801	San Diego, CA	24,899
Indianapolis, IN	7,055	Warren, MI	35,753
Sioux City, IA	7,223		

Table 2: Explaining variation in city pension expenditures per employee

	(1)	(2)	(3)
State collective bargaining law	0.333*** (0.073)		0.337*** (0.073)
Public-sector union membership		0.685*** (0.179)	
Ln(Population)	0.081** (0.034)	0.085** (0.034)	0.084** (0.034)
% Urban	0.316*** (0.121)	0.311*** (0.118)	0.318*** (0.121)
Ln(Income per capita)	0.186 (0.140)	0.187 (0.146)	0.185 (0.140)
Number of plans	0.007 (0.031)	-0.003 (0.032)	0.006 (0.031)
% Defined contribution plan	-0.233 (0.229)	-0.115 (0.228)	-0.236 (0.227)
% Local plan	-0.019 (0.132)	-0.031 (0.136)	-0.019 (0.132)
Includes OPEB	0.229* (0.118)	0.222* (0.116)	0.22* (0.118)
Issued POBs	1.166*** (0.348)	1.151*** (0.351)	1.24*** (0.345)
Constant	-1.308 (1.445)	-1.426 (1.494)	
Year fixed effects?	No	No	Yes
R-squared	0.24	0.23	0.27
Observations	1,745	1,745	1,745

Notes: Standard errors clustered by city in parentheses. Dependent variable is the log of total pension expenditures per employee per year. *p<0.1, **p<0.05, ***p<0.01.

Table 3: Effects of pension cost increases on city employment

	<i>Ln(Full-time equivalent employment per capita)</i>		<i>Ln(Full-time employees per capita)</i>		<i>Ln(Part-time employees per capita)</i>	
	(1)	(2)	(3)	(4)	(5)	(6)
Ln(Pension expenditures per employee)	-0.073*** (0.024)	-0.043** (0.021)	-0.077*** (0.023)	-0.042** (0.020)	-0.06 (0.058)	-0.049 (0.065)
Ln(General revenue per capita)	0.139** (0.057)	0.138** (0.064)	0.143** (0.058)	0.143** (0.064)	-0.004 (0.154)	-0.006 (0.155)
Issued POBs	0.096** (0.040)	0.049 (0.032)	0.103*** (0.037)	0.048 (0.029)	0.07 (0.171)	0.055 (0.172)
City fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	No	Yes	No	Yes	No	Yes
R-squared	0.98	0.98	0.98	0.98	0.91	0.91
Observations	1,461	1,461	1,461	1,461	1,445	1,445

Notes: Standard errors clustered by city in parentheses. All independent variables are lagged by one year. *p<0.1, **p<0.05, ***p<0.01.

Table 4: Collective bargaining and the effect of rising pension costs on full-time city employment

	(1)	(2)	(3)	(4)
Ln(Pension expenditures per employee)	-0.031 (0.021)	0.011 (0.023)	0.046 (0.058)	0.079 (0.064)
Collective bargaining *	-0.068** (0.034)	-0.077** (0.034)	-0.172 (0.121)	-0.187 (0.123)
Ln(Pension expenditures per employee)				
Ln(General revenue per capita)	0.143** (0.056)	0.144** (0.063)	0.136*** (0.052)	0.135** (0.059)
Issued POBs	0.09** (0.038)	0.032 (0.031)	0.111** (0.044)	0.052 (0.034)
Urban * Ln(Pension expenditures per employee)			-0.107* (0.061)	-0.095 (0.066)
Collective bargaining * Urban *			0.139 (0.127)	0.147 (0.126)
Ln(Pension expenditures per employee)				
City fixed effects?	Yes	Yes	Yes	Yes
Year fixed effects?	No	Yes	No	Yes
R-squared	0.98	0.99	0.98	0.99
Observations	1,461	1,461	1,461	1,461
Effect of pensions in bargaining states	-0.099*** (0.031)	-0.066** (0.027)		
Effect of pensions in rural cities, bargaining states			-0.126 (0.107)	-0.108 (0.101)
Effect of pensions in urban cities, non-bargaining states			-0.061*** (0.019)	-0.016 (0.019)
Effect of pensions in urban cities, bargaining states			-0.094*** (0.021)	-0.056*** (0.021)

Notes: Standard errors clustered by city in parentheses. Dependent variable is the log of the number of full-time city employees per capita. All independent variables are lagged by one year. *p<0.1, **p<0.05, ***p<0.01.

Table 5: Pension expenditures and employment reductions, by employee type

	<i>Ln(Fire protection employees per capita)</i>		<i>Ln(Police protection employees per capita)</i>		<i>Ln(Non-safety employees per capita)</i>	
	(1)	(2)	(3)	(4)	(5)	(6)
Ln(Pension expenditures per employee)	0.016 (0.033)	0.004 (0.035)	-0.017 (0.026)	0.017 (0.028)	-0.046* (0.027)	0.011 (0.029)
Collective bargaining *	-0.03 (0.047)	-0.028 (0.048)	-0.041 (0.032)	-0.047 (0.032)	-0.11** (0.044)	-0.123*** (0.043)
Ln(Pension expenditures per employee)						
Ln(General revenue per capita)	0.087 (0.068)	0.074 (0.066)	0.15*** (0.040)	0.143*** (0.043)	0.149** (0.060)	0.16** (0.065)
Issued POBs	0.035 (0.043)	0.053 (0.043)	0.033 (0.036)	-0.016 (0.034)	0.133*** (0.051)	0.056 (0.041)
City fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects?	No	Yes	No	Yes	No	Yes
R-squared	0.97	0.97	0.96	0.96	0.98	0.98
Observations	1,295	1,295	1,450	1,450	1,461	1,461
Effect of pensions in bargaining states	-0.014 (0.041)	-0.023 (0.040)	-0.058*** (0.021)	-0.030 (0.021)	-0.156*** (0.039)	-0.112*** (0.034)

Notes: Standard errors clustered by city in parentheses. All independent variables are lagged by one year. *p<0.1, **p<0.05, ***p<0.01.

Table 6: Pensions and capital outlays

	(1)	(2)
Ln(Pension expenditures per employee)	-0.225 (0.172)	-0.124 (0.168)
Collective bargaining *	-0.074 (0.184)	-0.081 (0.184)
Ln(Pension expenditures per employee)		
Ln(General revenue per capita)	0.773*** (0.228)	0.741*** (0.245)
Issued POBs	0.57*** (0.156)	0.409** (0.165)
Model	City fixed effects	City and year fixed effects
R-squared	0.65	0.66
Observations	1,505	1,505
Effect of pensions in bargaining states	-0.299*** (0.097)	-0.205* (0.113)

Notes: Standard errors clustered by city in parentheses. Dependent variable is logged per capita capital outlays. All independent variables are lagged by one year. *p<0.1, **p<0.05, ***p<0.01.