

Public Employee Pensions and Municipal Insolvency

Sean Myers

The Wharton School, University of Pennsylvania, USA

First version received October 2022; Accepted April 2026 (Eds.)

This paper studies how municipal governments jointly manage spending, credit market borrowing, and a public employee pension system. I model governments as levered investors who must meet non-defaultable pension obligations and may value government spending more than citizens. I quantify the model using data on California cities, including a new record of *fiscal emergencies*, tax increases required to maintain essential city services. After the financial crisis depleted pension funds, cities engaged in excessive risk-taking: the fiscal emergency option encouraged gambling for resurrection that kept cities vulnerable to shocks well into the recovery. To correct this problem, a savings requirement works better than a restriction on risk-taking or a pension funding requirement. The policy experiments emphasize that effective policies need to target the combined pension and bond finances, as policies that only target one, such as a pension funding requirement, are undermined by endogenous changes to the other.

Key words: Public pensions, Government debt, Risk management

JEL codes: H75, H74, H63

Across the U.S., state and local governments promise large pension benefits to public employees and invest in risky pension funds. Pension liabilities owed to employees totaled \$7.1 trillion in 2024 and pension fund assets totaled \$5.4 trillion, with 60% of assets invested in equity.¹ By borrowing in bonds and safe pension benefits and investing in risky pension funds, these governments act like hedge funds with levered exposure to risky assets. After the 2008 stock market bust, many governments could no longer maintain basic public services such as clean water or sewage and declared fiscal emergencies to ask local households for temporary tax increases. This begs the question, how do governments choose their risk exposure and should government policy be changed to reduce the risk of fiscal emergencies?

I propose a new model of government saving/borrowing which adds a second dimension to the government's problem: the choice of risk exposure. In this quantitative model, governments choose not only their bond saving or borrowing, but also their pension liabilities and investment in a pension fund that earns a risk premium. Each government provides public services for local

1. Center for Retirement Research at Boston College, Public Plans Data.

households and, when services are sufficiently low, the government can declare a fiscal emergency in order to ask households for a transfer. The model is quantified using a novel dataset of California city finances and fiscal emergencies, and I show that the model can explain both the saving and risk-taking behaviour of governments and has direct implications for government policy. Specifically, I find sizable welfare losses from governments being overly vulnerable to stock market busts. In two policy experiments, I find that savings requirements can address this issue and increase household welfare by 0.79% of consumption, whereas forcing the pension fund to invest in safer securities actually reduces household welfare by up to 0.39% of consumption. The policy experiments highlight the importance of jointly modelling the savings and risk exposure choices of the government, as policies which target one decision produce endogenous changes in the other which can negate the welfare gains for households.

To motivate the model, I document three stylized facts regarding municipal pensions and risk exposure. First, using settlement details for all U.S. city bankruptcies since 2008 as well as historical data for California back to 1967, I show that pension liabilities are senior to bonds. In California, they are effectively non-defaultable, meaning they are almost always fully paid. In other words, cities cannot easily walk away from these liabilities if they experience large losses on their pension funds. Second, using a panel of bond and pension finances for all California cities, I show that cities have high levered risk exposure. Across cities, both pension assets and pension liabilities are an order of magnitude larger than bonds outstanding. As a result, cities' balance sheets are mainly characterized by large non-defaultable liabilities and large risky assets which are mainly invested in equity. Third, the most indebted governments, measured as pension assets minus pension and bond liabilities, also have the most risk exposure. More indebted governments have larger safe liabilities and larger risky assets, meaning they are more exposed to stock market shocks.

I measure fiscal health using a hand-collected record of fiscal emergencies across all California cities. Fiscal emergencies (FE) are formal declarations that the government cannot maintain essential public services, and declaring a FE allows local households to vote on temporary tax increases. Because pension liabilities are non-defaultable and defaulting on bonds is costly, many cities declare a FE rather than defaulting. For governments that mainly have pension liabilities rather than bond liabilities, the cost of default outweighs the benefit of reducing their small bond liabilities.² After the financial crisis, 12% of California cities declared FE's while only 0.4% defaulted.

To understand how these three facts affect households and fiscal emergency risk, the model features governments that have the option to declare a FE if they cannot provide a threshold level of public services. This model builds on the sovereign default literature, which focuses on the savings decisions of governments who have outside options (*e.g.* Alfaro and Kanczuk, 2009; Chatterjee and Eyigungor, 2012; Bianchi *et al.*, 2018), by allowing the government to choose its risk exposure as well as its savings. In the model, government revenues and pension fund returns both depend on an aggregate shock. Each period, the government chooses how much to spend on public services, invest in the pension fund, and borrow or save in one-period bonds. Pension liabilities are non-tradable and accrue based on the amount the government spends on services. While the government can default on both pension liabilities and outstanding bonds, it endogenously chooses to only ever default on outstanding bonds, since public employees can punish the government much more severely than creditors. The decision to declare a FE in

2. The mayor of San Diego stated this exact rationale for not defaulting in 2010, "In the end, bankruptcy would cost hundreds of millions of dollars and net nothing in return. And the most compelling claim made on its behalf—that it would allow us to shed our pension obligations—is patently false."

order to ask households for a transfer is the result of a game between the government and local households. In equilibrium, the government declares a FE if services fall below a threshold. Public services cannot be traded, so the government's decisions cannot be offset through private household saving or borrowing.

A key element of the model is an agency conflict between households and the government. Consistent with the empirical evidence of risk shifting in public pension funds (Mohan and Zhang, 2014; Bradley *et al.*, 2016; Andonov *et al.*, 2017), the model government does not fully value the cost to households of providing transfers in a FE. This impacts the government's choice of both savings and risk exposure. Compared to a benevolent government that maximizes household utility, the model government saves less and takes more risk because it is partly insured by households against negative shocks. The effect intensifies as the government gets closer to a FE and transfers become more likely, which allows the model to explain why high debt governments take more risk in the data. As further support for this gambling for resurrection incentive, I document that prior to their bankruptcies, Detroit, Stockton, and San Bernardino all issued large amounts of bonds and placed the proceeds into their risky pension funds. In other words, these cities made substantial bets on the stock market prior to their bankruptcies.

Since model undersaving and excess risk become worse as a city approaches a FE, I quantify the model to match the response of cities to an unexpected drop in pension fund assets during the financial crisis. In order to compare cities that were more and less affected by the pension fund losses, I show that large cities tend to run bigger pension systems relative to revenue than small cities. This is due to large and small cities providing different mixes of services, with large cities spending more on labour-intensive services. Splitting cities into two groups based on population, this parsimonious model successfully matches a number of empirical moments. The model matches both the level and the year-by-year change in bonds, pension liabilities, and pension assets for both large and small cities after the financial crisis, as well as the total number of cities that declare a FE, the number of cities that default, and the average bond spread.³

Additionally, the model correctly predicts that large cities will be more likely to declare a FE, even though small cities receive more volatile revenue shocks.

With the panel of city finances and the quantified model in hand, I use counterfactuals and policy experiments to address the central questions of how vulnerable cities' saving and risk exposure decisions make them to stock market busts and whether changes in policy can improve household welfare. First, I test the ability of cities to handle another stock market bust similar to 2008.⁴ While one might expect that cities would increase savings and reduce risk exposure after the financial crisis, I find that the same shock to pension funds in 2015 would cause more than twice as many cities to declare a FE as the original 2008–10 shock: 136 out of 475 cities compared to only 57. This is because cities have chosen low savings and higher risk exposure since the financial crisis. While the panel ends in 2015, I also show that the positive returns on the pension fund for 2016–22 do not change this result.

Second, I calculate a constrained optimum where a benevolent government chooses spending and risk exposure. Compared to cities' actual choices after the 2008–10 shock, a benevolent government would choose to save substantially more and would reduce levered risk exposure, particularly for the cities closest to a FE. Starting at the 2010 empirical state variables, switching to the constrained optimal choices increases household welfare by 0.93% of consumption.

3. Trzcinka (1982), Yawitz *et al.* (1985), Stock (1994), Liu *et al.* (2003) and Schwert (2017) show that the large tax-adjusted spread on municipal bonds is explained by default premia, despite the low frequency of default. The model successfully matches this fact, generating both a low frequency of default and a high default risk premium.

4. Stock markets did experience a decline in early 2020, but this decline was quickly reversed within a few months and annual returns for stocks and pension funds were actually positive for both 2020 and 2021.

Given that the government undersaves and takes excessive risk, the policy experiments examine savings requirements and reducing the riskiness of the pension fund.⁵ A savings requirement is similar to a pension funding requirement but also includes bond borrowing or saving. While a city can meet a pension funding requirement by issuing more bonds and increasing pension fund assets, this will not change its overall savings. A savings requirement that limits spending to a constant, which is approximately 90% of current government spending, plus 0.4% of cash on hand increases household welfare by 0.79% of consumption. Forcing cities to save more moves them away from the FE region. As the likelihood of a FE decreases, government incentives become less skewed. As a result, the savings requirement also reduces the incentive to take excess risk and produces 85% of the welfare gain from moving to the benevolent government. Changing this policy from a savings requirement to a pension funding requirement almost entirely eliminates the welfare gain, as it primarily encourages the government to issue more bonds rather than to reduce spending.

In contrast to the savings requirement, requiring the pension fund to invest more in safe assets reduces household welfare.⁶ Forcing the pension fund to invest in safer assets lowers the average return of the fund, which reduces the government's incentive to save.⁷ Thus, it reduces excess risk-taking but worsens undersaving. In the quantified model, the latter force dominates and the result is a welfare loss for households. Investing the pension fund entirely in risk-free bonds lowers household welfare by 0.39% of consumption. Additionally, I find that household welfare monotonically decreases as the pension fund's portfolio weight on risk-free bonds increases. This means there is no intermediate combination of the current pension fund and a risk-free fund that produces a welfare gain for households.

Putting the counterfactuals and policy experiments together, the results emphasize two key points. The first is that effective policies need to target the city's combined pension finances and bond finances. Policies which only target one can be circumvented by making changes to the other. This is best highlighted by the ineffectiveness of the pension funding requirement compared to the savings requirement.⁸ The second is the importance of modelling savings and risk-taking as a joint decision by the government. Similar to the point about pension and bond finances, attempts to alter either the government's saving or risk-taking decisions will endogenously change its optimal choice for the other. In the case of the savings requirement, excess risk-taking is endogenously reduced as the government moves further away from the FE region which enhances the welfare gain from this policy. In the case of the restriction on risk-taking, governments endogenously reduce their savings which negates any welfare gains from reducing excess risk and actually leads to a welfare loss. Without a quantitative model of the joint saving and risk-taking decisions of the government, these endogenous responses will be overlooked.

Related literature. This paper builds on the large sovereign default literature (Eaton and Gersovitz, 1981; Aguiar and Gopinath, 2006; Arellano, 2008; see Aguiar *et al.*, 2016 for a survey) studying the savings decisions of governments who can relax their budget constraints in times of need and whose financial decisions endogenously impact the probability of using these

5. Because these are restrictions on the city government's choices, households would have to impose these policies through a state or county law, similar to existing restrictions on city government taxes. The city government would not voluntarily restrict its own choices.

6. As an example of this type of policy, until 1984 California only allowed pension funds to place 25% of their assets in stocks. Funds currently place roughly 60% in stocks.

7. Because city governments can default on their bonds, they cannot offset the change in the riskiness of the pension fund by simply issuing risk-free bonds à la Modigliani and Miller (1958). Any bonds the city issues would carry a default premium which would endogenously increase as the city issued more bonds.

8. Appendix E similarly shows that shifting from a savings requirement to a policy that only restricts bond issuance also eliminates all of the welfare gains.

outside options. A recent development in this literature is the inclusion of multiple tiers of liabilities. [Hatchondo *et al.* \(2017\)](#) study the potential creation of non-defaultable Eurobonds that are exogenously capped, and [Boz \(2011\)](#), [Fink and Scholl \(2016\)](#), and [Pancrazi *et al.* \(2020\)](#) analyse non-defaultable loans from institutions such as the IMF that carry “conditionality clauses” that restrict future government spending. This paper studies the choice of a government to issue defaultable and non-defaultable debt without caps or restrictions on spending.

A key feature of the model is that I allow the government to save in assets that earn a risk premium. [Alfaro and Kanczuk \(2009\)](#) and [Bianchi *et al.* \(2018\)](#) study the choice of a government to save in risk-free reserves to offset its defaultable bond liabilities. I also allow the government to save in a risk-free asset but extend its investment options to include risky assets that earn higher returns that are correlated with the government’s revenue growth. Because the government can default, it is partly insured against negative shocks, which increases its incentive to take risk. Thus, the core mechanism that drives borrowing and default in these models will also affect the incentive to take risk in my model. This incentive is strengthened by the fact that the government can also receive transfers from households in bad states through fiscal emergencies.

[Arellano *et al.* \(2016\)](#) and [Gordon and Guerron-Quintana \(2024\)](#) extend these sovereign default models to the sub-national level. While sub-national governments are often more constrained than federal governments due to legal restrictions and limited control over households, they are still qualitatively different from private companies. While private companies can be dissolved to pay off creditors, state and local governments endure which means their decisions will depend heavily on the continuation values. Both papers focus on the government’s choice of bonds, while my paper also incorporates the government’s choice of pension benefits and pension fund assets, which I show are an order of magnitude larger than bonds outstanding for California cities. Moreover, I find that 70% of bonds outstanding eliminated in recent city bankruptcies were bonds issued to invest in pension funds, so city pension decisions will be important for explaining bond borrowing and default.

This paper is also closely related to the empirical and theoretical literature studying unfunded public employee pensions.⁹ [Novy-Marx and Rauh \(2009, 2011\)](#), [Brown and Wilcox \(2009\)](#), and [Brown and Pennacchi \(2016\)](#) study the measurement of unfunded pensions and document their large size.¹⁰ Previous papers have separately modelled the choice of unfunded pensions and the choice of risk exposure for governments without outside options. The decision to underfund pensions has been studied by [Mumy \(1978\)](#), [Epple and Schipper \(1981\)](#), [Inman \(1982\)](#), and [Brinkman *et al.* \(2018\)](#) in deterministic economies and by [Glaeser and Ponzetto \(2014\)](#) when voter decisions are uncertain. [Lucas and Zeldes \(2009\)](#) and [Pennacchi and Rastad \(2011\)](#) study the government’s decision to invest its pension fund in safe or risky assets given an exogenous process for benefits and contributions. This paper combines the choice of underfunding and risk exposure and shows important interactions between the two decisions. Large unfunded pensions increase the chance of a fiscal emergency, which incentivizes governments to increase risk exposure, and restrictions on the riskiness of the pension fund affect the government’s incentive to fund pension benefits.

9. Unfunded pensions are the difference between pension liabilities and pension assets.

10. These papers calculate the liabilities for pension plans by discounting benefits at the return on safe assets and show that this is much larger than the values reported by governments, who discount at the average return on the pension fund. To account for this, the model focuses on a stream of benefit payments which is consistent with both sets of values. When discounted at the risk-free rate, this stream of payments matches the present values reported in these papers and when discounted at the average return on the pension fund, this stream of payments matches the present values reported by governments.

Layout. The remainder of the paper is organized as follows. Section 1 uses data on California cities to establish key stylized facts about savings, borrowing, and levered risk exposure and tests the defaultability of pension benefits. Section 2 describes a quantitative model in which the government can borrow or save in defaultable bonds, borrow in non-defaultable pension benefits, and save in a risky pension fund. Additionally, the government plays a game with households for fiscal emergency transfers. Section 3 quantifies the model to match the response of cities to the financial crisis. Section 4 estimates the risk of future fiscal emergencies, calculates the constrained optimum with a benevolent government, and examines counterfactuals and policy experiments. Section 5 shows that private savings by households cannot undo overspending by the government on public services and discusses why Ricardian equivalence does not hold. Section 6 concludes.

1. DATA AND FACTS

To understand how cities choose their risk exposure and borrowing and how this affects the chance of a fiscal emergency, I document several stylized facts about California city finances and create a novel dataset of fiscal emergencies.

Section 1.1 gives an overview of the pension system in California cities. Section 1.2 then presents a panel dataset of pension assets, pension liabilities, and bonds outstanding across all California cities for 2008–15. This panel shows that cities have high levered exposure to stock market risk and that the most indebted governments tend to have the highest risk exposure. Averaging across cities, pension assets and liabilities both dwarf the size of bonds outstanding in all years. Thus, cities are not simply entities that issue bonds to finance services but instead are more like hedge funds with large safe liabilities and large investments in risky assets. Across cities, there is considerable heterogeneity in the size of their balance sheets relative to revenue, with more indebted cities having larger balance sheets and greater risk exposure.

Section 1.3 then uses data on settlement details for all U.S. city bankruptcies since 2008 as well as historical data for California cities since 1967 to establish that pension liabilities appear to be senior to bonds and are effectively non-defaultable in California. Section 1.4 details a novel dataset of fiscal emergencies that I hand-collect for all California cities in order to test model predictions, as well as data on municipal bond spreads that will be used to estimate borrowing costs and default risk. Lastly, Section 1.5 discusses the use of discount rates when calculating pension liabilities and how to account for this by modelling a stream of benefit payments rather than a present discounted value.

1.1. *Pension system details*

City governments pay their employees with a mix of wages and defined pension benefits. For example, a retiring worker will know that based on her years of employment and salary history, she will receive \$50,000 a year plus cost of living adjustments for the rest of her life. To pay these future benefits, cities put money into a pension fund. These pension funds make risky investments; on average funds invest 60% in equity, 20% in bonds, and 20% in real estate and other assets. The majority of these pension funds are managed by the California Public Employees Retirement System (CalPERS), which acts as a third-party and invests all funds into the same portfolio. In terms of the average dollar, CalPERS manages 59% (63%) of all city pension funds weighted by pension fund assets (liabilities). In terms of the average city, CalPERS manages

95% of city pension funds.¹¹ The remaining 5% of pension funds choose a portfolio mix similar to CalPERS. As a consequence, the returns on pension funds vary from year to year, but do not vary across cities.

When measuring a city's pension finances, the main variables are the pension assets and liabilities. The pension assets are simply the amount of money in the city's pension fund. Legally, cities are not allowed to invest in risky securities except through their pension fund. The city government does not control the portfolio weights in the pension fund directly. Instead, it delegates the portfolio choice to third parties such as CalPERS. This means that the city government can only control its holdings of risky securities by putting more or less money into the fund. For example, some cities issue bonds and place the proceeds into the pension fund to increase their exposure to risk premia.

The pension liabilities reported in the data are the present value of benefits discounted at the average rate of return on the pension fund. In other words, the liabilities measure how much money needs to be in the pension fund today so that, on average, the pension fund will be able to pay these benefits in the future. In Section 1.5, I discuss alternative forms of discounting benefits and how the model addresses the use of different discount rates. Rather than including all benefits that employees are expected to earn over their careers, I only include a portion of the future benefits based on how much they have already worked.¹² This approach views pension benefits as a special type of bond and I only want to count the bonds that have already been used to pay employees.

1.2. *Panel of city finances*

I construct a panel of annual data on city finances from 2008 to 2015 for all California cities.¹³ The data merges information on city public pensions and their financial statements. The pension data is taken from actuarial valuation reports provided by CalPERS or the individual cities.¹⁴ Most city pension plans are run through CalPERS, which manages the investments of the funds and payouts to retirees. For plans that are not run through CalPERS, the data is collected from the Required Supplementary Information section of their Comprehensive Annual Financial Report. The main variables are the pension assets and pension liabilities.

Data on city bonds outstanding and revenue comes from the City Financial Transactions Reports (CFTR). Each year, cities must file this report with the State Controller's Office (SCO) based on their audited financial statements. For the small number of instances in which a city did not file their CFTR with the SCO by the reporting deadline, I manually collect their audited financial statements.¹⁵ The main variables from these reports are annual city revenues and bonds outstanding. Since many small cities mainly use bank loans instead of issuing municipal bonds, I also include bank loan debt in the measure of bonds outstanding.

11. This distinction between the average across cities and the asset-weighted or liability-weighted average comes from the fact that the cities that manage their own pension funds tend to be the largest cities in the state.

12. The liabilities are calculated using the Entry Age Normal actuarial method. Under this method, the accrued liability owed to employees increases by a constant percentage of their wages each year that they work. This aligns well with the model, where new liabilities accrue as a constant percentage of total expenditure on services.

13. The panel ends in 2015 as this is the final year for the hand-collected fiscal emergency records discussed in Section 1.4.

14. I am grateful to Joe Nation and Pension Tracker for providing access to the pension data.

15. Three cities (Loyalton, Tulalake, Westmorland) are excluded from the panel due to missing data. Combined these cities only represent 0.012% of the total population for the 478 California cities. For Tulalake and Westmorland, the underlying audit could not be found for one of the missing years. For Loyalton, the city did not report pension liabilities and assets for 2013–2015 because of a lawsuit with employees.




FIGURE 1
Time series of city finances

Notes: This figure shows the time-series for average California city pension liabilities (in dark red), pension assets (in dashed blue), and bonds outstanding (in light yellow) as a percent of revenue. For each (city, year) pair, pension liabilities, pension assets, and bonds outstanding are measured as a percent of revenue. Then, these values are averaged across cities to produce the time series.

In Appendix B, I compare the aggregate and city-by-city values for my bonds outstanding measure to the “total outstanding debt” measure from the Census of Governments for the years in which the census is conducted. As one might expect, bonds outstanding tends to be a lower bound on total debt outstanding, as total debt outstanding includes additional items such as funds owed to the state or federal government. In 2012, which is the latest year in the panel for which the census is conducted for all cities, the aggregate value of total debt outstanding across all cities is \$95.4 billion and the aggregate value of bonds outstanding is roughly half that amount at \$44.2 billion. I focus on bonds outstanding for my analysis so that the data most closely aligns with the model environment of privately traded claims issued by the city government which carry a market-based risk premium.

1.2.1. Time series. Figure 1 shows the time series for pension liabilities, pension assets, and bonds outstanding as a percent of revenue averaged across all cities. Figure A1 shows the same series using the population-weighted averages across cities. The most prominent feature of this graph is that the pension system dwarfs bonds outstanding. The cross-sectional average of pension liabilities and assets relative to revenue is always above 100%, while bonds outstanding relative to revenue are about 20%.

Because the pension system is characterized by investment in a risky pension fund to cover promises for safe pension benefits, the large size of the pension system means that cities have levered risk exposure. Put succinctly, cities have large risky assets and large safe liabilities. Solely focusing on the difference between pension liabilities and assets, *i.e.* the unfunded pensions, misses this risk exposure. For example, having liabilities of 10% revenue and assets of 0% revenue would give the same value for unfunded pensions as having liabilities of 110% revenue and assets of 100% revenue. However, in the latter scenario, the city is much more exposed to stock market risk and unfunded pensions could double or completely disappear with just a 10% unexpected loss or gain on the risky assets.

The financial crisis highlights this large exposure to risk. In 2008, pension liabilities and pension assets averaged 149% and 139% of revenue, respectively. Unfunded pensions were thus 10% of revenue and were smaller than bonds outstanding, which averaged 19% of revenue. However, pension assets dropped during the financial crisis whereas the safe pension liabilities continued to grow, causing unfunded pensions to sextuple to 64% of revenue in 2009.¹⁶ While the pension fund has delivered positive returns since 2009, the cumulative return has not caught up to its pre-crisis trend, meaning that pension assets are still substantially below the value cities were forecasting in 2008. Even in the final panel year of 2015, the gap between pension liabilities and pension assets remains large at 55% of revenue.¹⁷

Importantly, the financial crisis is not the sole reason for the current state of unfunded pensions in California cities. For example, after unexpected positive shocks to pension fund assets in the 1990's, many cities reduced or suspended contributions to their pension funds and, at the state level, California retroactively increased existing pension benefits. These practices did not result in large unfunded pensions at the beginning of the panel in 2008. But they did ensure that pensions were not overfunded and were vulnerable to a large stock market decline, and the financial crisis was the realization of one such decline. To capture these interactions, the model of Section 2 features cities that choose their pension contributions—including the possibility of making no contributions—in an environment with risky pension assets.

1.2.2. Cross-section. There is substantial heterogeneity in debt and risk exposure across cities. Figure 2 shows the cross-section of city finances in 2015, the final year of the panel. There are three important differences to note. First, there is wide variation in the size of city balance sheets, with many cities having both assets and total liabilities below 50% of revenue, while others have assets and total liabilities in excess of 300% revenue. Second, cities differ substantially in their debt, as shown by the net position dots of Figure 2. Debt is the amount that pension and bond liabilities exceed pension assets, or equivalently a city's bonds plus unfunded pensions. While debt as a percent of revenue is 71% on average, the black dots of Figure 2 show that this ranges from virtually 0% to immense values above 300%. These two results both indicate that there are some cities which are much more vulnerable to a negative stock market shock. Towards the left-side of the graph, there are many cities that have almost no debt and are fairly insulated from any financial shocks because they do not have large assets or total liabilities. Towards the right-side of the graph, we have many cities with substantial debt that also are highly exposed to financial shocks because they have a levered position in the risky assets, with large pension assets and large total liabilities.

The third difference to note is that the composition of total liabilities varies substantially across cities. While pension liabilities are on average much higher than bonds outstanding at 217% of revenue compared to only 16.4%, bonds constitute a large portion of total liabilities for some cities. For 20 cities, bonds outstanding are actually larger than pension liabilities. This means that any difference in the treatment of bonds and pension liabilities will matter when assessing which cities are more likely to declare fiscal emergencies. For example, if pension liabilities can be easily renegotiated, then we may not need to worry about cities where total liabilities are predominantly pension liabilities. Section 1.3 shows that pension liabilities actually appear to be senior to bonds in recent bankruptcies and effectively non-defaultable in California,

16. This same pattern holds across the U.S. The national aggregate of state and local public employee pensions from the Financial Accounts of the United States shows that pension assets fall during the financial crisis, creating a persistent gap between liabilities and assets.

17. Section 4.1 shows that positive returns for the pension fund from 2016 to 2022 do not close this gap. The pension fund's cumulative return is still noticeably below its pre-crisis trend.




FIGURE 2

Cross-section of 2015 pension assets, pension liabilities, and bonds outstanding

Notes: This figure shows the pension assets, pension liabilities, and bonds outstanding as a percent of revenue for each of the 475 California cities in 2015. Cities are sorted by their total pension and bond liabilities as a percent of revenue. The blue positive bars represent the pension assets, the dark red negative bars represent the pension liabilities, and the light yellow negative bars represent bonds outstanding. The black dots show the net position of the city, which is the pension assets minus the pension liabilities and bonds outstanding.

meaning that large pension liabilities are more likely to drive a city into a fiscal emergency than large bonds outstanding.

1.2.3. Large and small city differences. Combining Figures 1 and 2, we know that debt increased significantly during the financial crisis, as pension assets fell while pension liabilities continued to grow, and that this did not affect all city balance sheets equally, as some cities run much bigger pension systems relative to revenue than others. For the quantification exercise of Section 3, it will be useful to separate cities that were more and less affected by the increase in debt. One of the factors that contributes to the large heterogeneity across cities in their balance sheets is the fact that cities do not all provide the same bundle of services. Large population cities spend more of their revenue on public safety, which primarily uses pensioned labour in the form of police and firefighters. In comparison, small population cities spend more on services such as sewage, which are more capital intensive and use non-pensioned labour. As a consequence, large cities have higher pension liabilities relative to revenue than small cities. Large cities can cover these liabilities by issuing fewer bonds or saving more in their risky pension funds.

To focus on these differences, I sort cities into two groups based on their 2010 Census populations. For both groups, I measure the payroll for pensioned employees as well as the expenditure on public safety and public health. This data comes from the CalPERS valuation reports, the Comprehensive Annual Financial Reports, and the CFTR's. Large cities spend 25.0% (0.2%) of revenue on wages for pensioned labour, compared to only 20.6% (0.3%) for small cities, with the standard errors reported in parentheses. This 4.4 (0.4) percentage point difference is almost exactly explained by differences in public safety and public health expenditures. Large cities spend 4.5 (0.4) percentage points more of revenue on public safety than small cities and 4.5 (0.5) percentage points less of revenue on public health, which is primarily sewage and waste.

1.3. *Seniority of pensions*

Do cities always pay their pension liabilities? Given the large size of pension liabilities, it is important to know if they are state-contingent and can be defaulted on and cut after bad shocks. This section provides evidence that pension liabilities are senior to bonds, meaning they are not substantially cut even in bankruptcy. For California in particular, pension liabilities appear to be virtually non-defaultable, with almost no instances of accrued benefits not being fully paid.

To check whether pension liabilities can be easily altered, I collect settlement details for all five city bankruptcies across the U.S. since 2008.¹⁸ Cuts to bonds are taken from Moody's and cuts to pension liabilities are taken from the bankruptcy exit plans. Because pension liabilities include all accrued benefits, liabilities will still be cut even if a city leaves benefits for current retirees unchanged and simply alters the already accrued benefits for current workers. For the five bankruptcies, bonds outstanding were on average cut by 29% while pension liabilities were only cut by 10%. This is despite the fact that bonds outstanding averaged only 71% of revenue while pension liabilities averaged 352%. So, cities chose to make larger cuts to bonds even though they could have saved more money by cutting pensions.

In California, the results are even more stark. Over 95% of cities' pension plans are run through CalPERS and only one has ever failed to fully pay accrued benefits since CalPERS began managing city pensions in 1967. This one instance in Loyalton only affected three employees. In other words, pension liabilities have been close to non-defaultable for more than five decades. Interestingly, this is not simply due to legal protections for the pensions. In multiple cases, judges ruled that California cities could default on accrued benefits, but even cities in bankruptcy still chose to fully pay the pension liabilities.¹⁹ Cities have stated that the main reason they do not cut pensions is because public employees may move to other cities if the government breaks its promise with workers. During Stockton's bankruptcy, the city manager stated that defaulting on pension plan benefits would lead to a "mass exodus" of public employees. Similarly, a disclosure from the San Bernardino bankruptcy states "the City determined that it could not reject its contract with CalPERS. The City concluded that rejection of the CalPERS contract would lead to an exodus of City employees."

1.4. *Fiscal emergencies and bond spreads*

In the presence of non-defaultable pension liabilities, default risk may be a poor measure of a city's fiscal health. The reason is that a city may struggle to pay its large unfunded pensions but still choose not to default unless it also has large bond liabilities that will be reduced in bankruptcy. In 2010, the mayor of San Diego, Jerry Sanders, expressed this idea when he wrote "In the end, bankruptcy would cost hundreds of millions of dollars and net nothing in return. And the most compelling claim made on its behalf—that it would allow us to shed our pension obligations—is patently false."²⁰

The measure of fiscal health that I will study is fiscal emergencies. In California, tax changes require voter approval and this vote can only be held once every two years. A fiscal emergency is a declaration of service-level insolvency that allows a city to skip this waiting period and hold an immediate vote on taxes. The key requirement is that the city must show that this is necessary

18. These settlement agreements exist for cities that successfully enter bankruptcy and exclude cities that file for bankruptcy but have their application denied or withdrawn.

19. Boyer (2020) reaches a similar conclusion for the California state government. He infers from bond spreads that 90–100% of unfunded pensions would be untouched even after a 100% haircut to bonds.

20. Sanders, Jerry, "Debunking the bankruptcy myth", San Diego-Union Tribune (2010).

to avoid cutting essential services, such as water, sewage, or public safety. Unlike default, both bond liabilities and pension liabilities could potentially force a city to cut services and declare a fiscal emergency. Fiscal emergencies are also more directly tied to household welfare than default risk, as they relate specifically to the public services provided to residents.

To measure these events, I hand-collect a record of fiscal emergencies in California cities from 2000 to 2015. I create this record by contacting the clerk's office of all 475 cities and running a public records search for declarations of a fiscal emergency. As an example, in 2013 the city of Antioch declared a fiscal emergency in order to ask voters for approval of a temporary sales tax increase, designed to pay for essential services. The title section of Resolution No. 2013/33 summarizes the declaration:

Resolution of the city council of the City of Antioch declaring a Fiscal Emergency; calling for and noticing a municipal election on November 5, 2013 to present to voters a temporary one-half cent transactions and use (sales) tax to fund all essential Antioch city services including police, code enforcement, and economic development.

The resolution then details the recent cuts to public services and how this qualifies as being unable to provide "essential services." In this case, the main factors are a 31% reduction in police department staffing, a 30% increase in violent crime, a 23% increase in property crime, and an increase in response times for priority 1 (life threatening) police calls.

While bond spreads and defaults may not be as useful as fiscal emergencies for measuring fiscal health, they are still important for understanding city borrowing costs. A city's decision to issue bonds will depend heavily on the interest rate it faces. To discipline borrowing costs in the model, I collect data on all municipal bond trades where customers are buying or selling to a dealer from MSRB for 2008–15. I then identify the bonds associated with each city using the 6 character issuer CUSIP's from the Thomson Reuters MuniProfiles database and collect the characteristics for each bond from Bloomberg such as its tax status and whether the bond is insured by a third-party.

1.5. Discussion of discount rates

Cities report their liabilities as the value of benefit payments discounted at the average return on the pension fund.²¹ However, since the pension fund earns a risk premium and pension benefits are risk-free, this understates the value of these benefits. As argued by [Brown and Wilcox \(2009\)](#) and [Novy-Marx and Rauh \(2011\)](#), if pension benefits are safe claims then using a risk-free discount rate is the correct method to measure the present value of benefits. This gives substantially higher values for pension liabilities. For 2015, the liabilities of CalPERS city plans discounted at a risk-free rate of 3.5% were 1.65 times the value discounted at the standard 7.5%.

Another way to view this difference is that cities are not reporting the value of insuring the pension fund against risk. While the reported liabilities state how much the pension fund needs to cover benefits on average, the government is responsible for any losses or gains from the pension fund under- or over-performing. The value of this insurance is precisely the difference between the value of benefits discounted at the average return on the pension fund and the value of benefits discounted at the risk-free rate.

21. As noted in [Andonov et al. \(2017\)](#), if a government can control the riskiness of the pension fund then it could change its reported liabilities by changing the discount rate. For California cities, this is not an option as the city does not control the riskiness of the pension fund. Most city pension funds are managed by CalPERS which invests all cities' funds into the same portfolio, and the cities that do not operate through CalPERS delegate the management to other parties. Even if a city issues bonds and invests the proceeds into the pension fund, this will not change the discount rate as the return of the pension fund is unchanged.

To address this issue, the model of Section 2 uses a stream of benefit payments rather than a present discounted value. I approximate pension benefits as long-term bonds where a constant fraction mature each period following the standard formulation for long-term bonds of [Leland \(1994\)](#), [Hatchondo and Martinez \(2009\)](#), [Arellano and Ramanarayanan \(2012\)](#), and [Chatterjee and Eyigungor \(2012\)](#). I set the fraction of benefits that mature each period to match the empirical duration of pension benefits, *i.e.* the sensitivity of liabilities to changes in the discount rate.²² Matching the empirical duration ensures that the model stream of benefit payments is consistent with both the present value reported by cities and the present value calculated using a risk-free discount rate.

2. MODEL

To understand how city finances affect fiscal emergency risk, and conversely how the chance of a fiscal emergency affects the cities' financial decisions, I construct a dynamic, small open economy model. The model features two main agents. The first is a government that provides public services to households. The government pays for these services with a combination of wages and pension benefits. The government can borrow or save in bonds and save in a risky pension fund. It can also default on pension benefits or bonds outstanding but at different costs and can declare a fiscal emergency in order to ask households for a transfer. In equilibrium, the cost of defaulting on pensions will be high enough that it never occurs.

The second main agent is households, who have preferences over public services as well as their own consumption. For this section, I make the simplifying assumption that households do not borrow or save. Section 5 presents an extended model in which households can privately borrow and save and shows that the results do not noticeably change. In other words, having households that simply consume their income is a reasonable approximation of the outcome when households can privately borrow/save and significantly simplifies the computation of the equilibrium. If a fiscal emergency is declared, then households will choose how much they want to transfer to the government to pay for additional public services. In addition to the two main agents, the model will feature creditors who price all of the assets, including the municipal bonds.

2.1. Endowments, preferences, and assets

2.1.1. Governments. In each city, public services are provided by a local government. Each city government receives an exogenous endowment of tax revenue. The tax revenue is exogenous to capture the fact that California cities cannot independently change taxes. This revenue fluctuates based on i.i.d. aggregate growth shocks as well as an idiosyncratic log AR(1) component. Specifically, the government of city i has revenue

$$Y_{i,t} = \Gamma_t y_{i,t}. \quad (1)$$

The component Γ_t is shared across all cities and is subject to growth shocks g_t ,

$$\Gamma_t = g_t \Gamma_{t-1}. \quad (2)$$

22. Since 2013, CalPERS plans have reported their liabilities discounted at a risk-free rate as well as the average return on the fund. The sensitivity of liabilities to the change in the discount rate implies an average duration of 13 years.

The idiosyncratic component $y_{i,t}$ follows

$$\log(y_{i,t}) = \rho \log(y_{i,t-1}) + \sigma \varepsilon_{i,t} \quad (3)$$

where $\varepsilon_{i,t}$ has a standard normal distribution. The government has preferences

$$\mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t u(S_{i,t}) \right] \quad (4)$$

over public services $S_{i,t}$ with CRRA felicity, $u(S_{i,t}) = \frac{1}{1-\gamma} S_{i,t}^{1-\gamma}$.

2.1.2. Households. Each city i has local households. I make the simplifying assumption that these households do not borrow or save. Section 5 removes this assumption and shows that the quantitative predictions of the model are almost the same when households can privately borrow/save. These households have disposable income

$$Y_{i,t}^h = \Gamma_t y^h \quad (5)$$

which can either be consumed or given to the government to spend on services. The value for y^h determines the size of this disposable income relative to government revenue. Households have preferences

$$\mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t (u(S_{i,t}) + v(C_{i,t})) \right] \quad (6)$$

over public services $S_{i,t}$ and household consumption $C_{i,t}$. Households have the same felicity for services as the government, $u(S_{i,t})$, and have CRRA felicity for consumption, $v(C_{i,t}) = \psi \frac{1}{1-\gamma} C_{i,t}^{1-\gamma} - \psi \frac{1}{1-\gamma} (Y_{i,t}^h)^{1-\gamma}$. The felicities $u(\cdot)$ and $v(\cdot)$ have the same curvature γ to ensure a balanced growth path where the desire for consumption relative to services does not shrink to 0 or explode to ∞ as Γ_t grows. Because $Y_{i,t}^h$ is an entirely exogenous process, the $\psi \frac{1}{1-\gamma} (Y_{i,t}^h)^{1-\gamma}$ term in $v(C_{i,t})$ has no impact on household choices; it is solely included to simplify comparisons between government felicity and household felicity.

2.1.3. Creditors. Each individual city is a small open economy. Outside of the cities is a continuum of competitive creditors with consumption $\Gamma_t y^m$ and preferences

$$\mathbb{E} \left[\sum_{t=0}^{\infty} \beta_m^t v^m(\Gamma_t y^m) \right]. \quad (7)$$

Since the creditors determine the price of municipal bonds, it is useful to represent their preferences in terms of their pricing kernel

$$M_t = \beta_m \frac{v^{m'}(\Gamma_t y^m)}{v^{m'}(\Gamma_{t-1} y^m)}. \quad (8)$$

2.1.4. Assets. Each government's portfolio is comprised of three components. The first component is the pension fund assets $A_{i,t}$. The government can save in a risky pension fund that has return $R(g_t) > 0$ which varies depending on the aggregate shock. The return on the risky pension fund is identical across cities, as explained in Section 1.1. However, cities can still

take on more or less risk through the use of one-period bonds $B_{i,t}$ which are held by creditors. If the government saves in bonds, it earns a risk-free rate $R_t^f = 1/E[M_{t+1}]$. If the government borrows in bonds, then the bond price depends on the probability the government will default in the next period. By choosing how much to save or borrow in bonds and how much to save in the risky pension fund, cities can control their exposure to aggregate risk.

The final component is pension benefits $B_{i,t}^p$. These benefits are non-tradable; they are accrued when the government provides services and disbursed when benefits mature. The cost of services is normalized so that one unit of services costs the government one unit in immediate wages and η units in pension benefits. Each period, a portion λ of these pension benefits mature, meaning that the government has a stream of geometrically decaying payments $\lambda, (1 - \lambda)\lambda, (1 - \lambda)^2\lambda, \dots$ for each unit of $B_{i,t}^p$.

2.2. Default decision

At the beginning of each period, the government either pays or defaults on maturing pension benefits and bonds outstanding. If the government does not pay maturing pension benefits $\lambda B_{i,t}^p$, then services $S_{i,t}$ are zero. This reflects the finding in Section 1.3 that governments fear a “mass exodus” of employees if they default on pension benefits. Because $u'(0)$ is infinite, this punishment is strong enough that pensions are always paid.

If the government defaults on bonds $B_{i,t}$, then the creditors and the city go to court and a bankruptcy judge forces the city to pay a default cost of $\phi(g_t)Y_{i,t}$ to creditors. The default cost can vary with the aggregate state to allow for the possibility that default may be easier in bad aggregate states. This is discussed more in Section 3.1. While many sovereign default models also include a period of exclusion from credit markets as part of the punishment for default, I do not assume that cities are excluded after default. This is done to allow for the possibility that cities pay the default cost on their current bonds outstanding partly through issuing new bonds, as often occurs in city bankruptcies.²³ Given that payment $\phi(g_t)Y_{i,t}$ to creditors is the only punishment, default on bonds occurs whenever $B_{i,t} > \phi(g_t)Y_{i,t}$, i.e. the cost of paying the bonds exceeds the default cost. The competitive bond price is then the discounted value of future payments,

$$q(B_{i,t+1}, \Gamma_t, y_{i,t}) = \mathbb{E} \left[M_{t+1} \mathbb{1} \{ B_{i,t+1} \leq \phi(g_{t+1}) Y_{i,t+1} \} \right] + \mathbb{E} \left[M_{t+1} \mathbb{1} \{ B_{i,t+1} > \phi(g_{t+1}) Y_{i,t+1} \} \frac{\phi(g_{t+1}) Y_{i,t+1}}{B_{i,t+1}} \right]. \quad (9)$$

Because the future payments only depend on $B_{i,t+1}$ and $\phi(g_{t+1})Y_{i,t+1}$, we have a simple bond price schedule that only depends on the chosen $B_{i,t+1}$ and the variables that impact future $Y_{i,t+1}$ (i.e. Γ_t and $y_{i,t}$).

2.3. Game between government and households

I focus on the Markov Perfect Equilibrium of a sequential stage game between households and the government. From here forward, I drop the i and t subscripts.

23. The lack of exclusion is also consistent with the assumption that creditors are competitive. If creditors attempt to collude and refuse to buy a city's new bonds after default, some creditors will always deviate and buy the bonds at a low price. In comparison, unionized public employees can much more credibly threaten to coordinate on punishing the government. This would continue until enough creditors deviate that the bonds are competitively priced.

The timing of the model is the following: at the beginning of each period, the exogenous states Γ, g, y are realized and bonds are defaulted on if $B > \phi(g)Y$. The city government then chooses whether to declare a fiscal emergency δ , spending on services S^g , and the portfolio $(A', B', B^{p'})$. If the government does not declare a fiscal emergency ($\delta = 0$), then the game is over for this period and public services and household consumption are simply S^g and Y^h . If the government does declare a fiscal emergency ($\delta = 1$), then households choose an amount to transfer $S^h \geq 0$ which will be taken from their consumption and spent on services. This is designed to capture that, during fiscal emergencies in the data, households are voting on a temporary tax increase to fund public services. In the model, households provide the government with a temporary transfer to increase services. Formally, we have

$$S = S^g + \delta S^h, \quad (10)$$

$$C = Y^h - \delta S^h. \quad (11)$$

2.3.1. Household decision. In the second stage, households choose a transfer S^h if a fiscal emergency was declared by the city government. Given that the government chooses the entire portfolio $(A', B', B^{p'})$ for the next period in the first stage, households cannot affect the future state variables. Thus, households choose the transfer to maximize current felicity $u(S) + v(C)$ subject to the constraints (10), (11), $S^h \geq 0$, and $B^{p'} \geq (1 - \lambda)B^p + \eta S$. This gives that the optimal transfer is

$$S^h(S^g, \Gamma, B^p, B^{p'}) = \min \left\{ \hat{S}^h(S^g, \Gamma), [B^{p'} - (1 - \lambda)B^p] / \eta - S^g \right\} \quad (12)$$

where the function $\hat{S}^h(S^g, \Gamma)$ is defined as

$$\hat{S}^h(S^g, \Gamma) = \max \left\{ \frac{1}{1 + \psi^{1/\gamma}} Y^h - \frac{\psi^{1/\gamma}}{1 + \psi^{1/\gamma}} S^g, 0 \right\}. \quad (13)$$

Note that when $\hat{S}^h(S^g, \Gamma)$ is strictly positive, its value $\frac{1}{1 + \psi^{1/\gamma}} Y^h - \frac{\psi^{1/\gamma}}{1 + \psi^{1/\gamma}} S^g$ corresponds to the interior solution of choosing S^h to maximize household felicity $u(S^g + S^h) + v(Y^h - S^h)$. In other words, the transfer $\frac{1}{1 + \psi^{1/\gamma}} Y^h - \frac{\psi^{1/\gamma}}{1 + \psi^{1/\gamma}} S^g$ equates the marginal felicity of services $u'(S^g + S^h)$ and the marginal felicity of consumption $v'(Y^h - S^h)$.

2.3.2. Government decisions. In the first stage, the government enters the period with exogenous states (Γ, y, g) and portfolio (A, B, B^p) . After paying maturing pension benefits and either paying or defaulting on bonds, the government has cash-on-hand

$$X = Y - \lambda B^p - \min \{B, \phi(g)Y\} + R(g)A, \quad (14)$$

and non-tradable pension benefits $(1 - \lambda)B^p$.

Given the household transfer function $S^h(S^g, \Gamma, B^p, B^{p'})$, the government's problem is

$$\begin{aligned}
 V(X, B^p, \Gamma, y) &= \max_{S, S^g, \delta, A', B', B^{p'}} u(S) + \beta \mathbb{E}[V(X', B^{p'}, \Gamma', y') | \Gamma, y] \\
 S^g &= X - A' + q(B', \Gamma, y) B' \\
 S &= S^g + \delta S^h(S^g, \Gamma, B^p, B^{p'}) \\
 B^{p'} &\geq (1 - \lambda) B^p + \eta S \\
 \delta S^h(S^g, \Gamma, B^p, B^{p'}) &\geq 0 \\
 A' &\geq 0 \\
 X' &= \Gamma' y' - \lambda B^{p'} - \min\{B', \phi(g') \Gamma' y'\} + R(g') A'.
 \end{aligned} \tag{15}$$

The budget constraint states that government spending on services is equal to cash-on-hand, minus investment in pension assets, plus (minus) any amount that the government borrows (saves) in bonds. Total services are simply the amount of government spending on services plus any fiscal emergency transfers from households. Next period pension benefits must be at least the non-matured benefits plus newly accrued benefits from services. Transfers from households must be nonnegative. And lastly, the government cannot borrow in the pension assets.

Conveniently, whenever the government declares a fiscal emergency ($\delta = 1$), we can show that it will choose S^g and $B^{p'}$ such that $S^h(S^g, \Gamma, B^p, B^{p'}) = \hat{S}^h(S^g, \Gamma)$. Intuitively, if it was ever the case that the government declared a fiscal emergency and $\hat{S}^h(S^g, \Gamma) > [B^{p'} - (1 - \lambda)B^p]/\eta - S^g$, then the government could hold $B^{p'}$ and B' fixed, slightly increase A' , and slightly decrease S^g without changing $S = S^g + \delta S^h(S^g, \Gamma, B^p, B^{p'})$. This change strictly increases X' which increases the continuation value. Because the government does not value that transfers reduce household private consumption, the government would always make this change until it satisfied $S^h(S^g, \Gamma, B^p, B^{p'}) = \hat{S}^h(S^g, \Gamma)$. It is also straightforward to show that the government will always choose $B^{p'} = (1 - \lambda)B^p + \eta S$ since the continuation value is strictly decreasing in $B^{p'}$.

If the government is indifferent between declaring and not declaring a fiscal emergency, I follow the literature on sovereign default and break the indifference by assuming that the government does not declare a fiscal emergency. This means that we can further simplify the problem by noting that the government declares a fiscal emergency if and only if $\hat{S}^h(S^g, \Gamma) > 0$. This is shown formally in the footnote.²⁴ The intuition is that because the government only cares about services, it will always accept any transfers that the households are willing to offer. This also maps nicely to the data where governments declare a fiscal emergency only if they need money for “essential services.” In the model, the level of essential services will be the threshold for spending S^g such that $u'(S^g)$ exceeds $v'(Y^h)$.

24. If $\hat{S}^h(S^g, \Gamma) = 0$, then the government is indifferent between $\delta = 0$ and $\delta = 1$. Given that I break the indifference by assuming the government does not declare a fiscal emergency when indifferent, this means $\delta = 0$. Thus, it only needs to be shown that $\delta = 1$ if $\hat{S}^h(S^g, \Gamma) > 0$. Suppose for some (X, B^p, Γ, y) the government chooses $\delta = 0$ and S^g, A' such that $\hat{S}^h(S^g, \Gamma) > 0$. Then there is a profitable deviation. Let $\tilde{Z}$ denote the deviation value of any variable Z . Starting at $(\tilde{\delta}, \tilde{S}^g, \tilde{A}') = (1, S^g, A')$, the government can lower $\tilde{S}^g$ and equally raise $\tilde{A}'$ until satisfying $\mathbb{E}[V(\tilde{X}', \tilde{B}^{p'}, \Gamma', y') | \Gamma, y] = \mathbb{E}[V(X', B^{p'}, \Gamma', y') | \Gamma, y]$. If $\tilde{A}' = A'$ at this point, then $\tilde{S}^g = S^g$, which means $\tilde{S} > S$. If $\tilde{A}' > A'$, then equality of the continuation values implies that $\tilde{B}^{p'} > B^{p'}$, which means $\tilde{S} > S$. In both cases, this deviation gives the same continuation value but higher present services. Therefore, the government will always choose $\delta = 1$ if $\hat{S}^h(S^g, \Gamma) > 0$.

Thus, the government's problem can be reduced to

$$\begin{aligned}
 V(X, B^p, \Gamma, y) &= \max_{S^g, A', S, B', B^{p'}} u(S) + \beta \mathbb{E}[V(X', B^{p'}, \Gamma', y') | \Gamma, y] \\
 S^g &= X - A' + q(B', \Gamma, y) B' \\
 S &= S^g + \hat{S}^h(S^g, \Gamma) \\
 B^{p'} &= (1 - \lambda) B^p + \eta S \\
 A' &\geq 0 \\
 X' &= \Gamma' y' - \lambda B^{p'} - \min\{B', \phi(g') \Gamma' y'\} + R(g') A'.
 \end{aligned} \tag{16}$$

Each period, the government just has to choose spending S^g and its risky pension assets A' . The values for B' , S , and $B^{p'}$ are then determined by the constraints of equation (16), namely the budget constraint, the fiscal emergency transfers, and the law of motion for pension benefits.²⁵ Higher S^g means the government is spending more today and saving less for the future. Higher A' means the government is saving more in risky assets and saving less or borrowing more in bonds.

2.4. Equilibrium

A Markov Perfect Equilibrium consists of a value function V , policy functions $\{S, S^g, A', B', B^{p'}\}$, and a transfer function $\hat{S}^h$ that satisfy the following conditions:

1. Taking as given the transfer function $\hat{S}^h$, the policy functions $\{S, S^g, A', B', B^{p'}\}$ and the value function V solve the government's problem (16).
2. The transfer function $\hat{S}^h$ satisfies (13).

2.5. Transfers skew incentives

A key aspect of the game between households and the government is that the ability of the government to receive transfers in bad states will skew its incentives. Using equation (13) and the decision rule for declaring a fiscal emergency from Section 2.3, $\delta = 1$ if and only if $\hat{S}^h(S^g, \Gamma) > 0$, we have that equilibrium services and consumption are

$$S(S^g, \Gamma) = \max \left\{ \frac{1}{1 + \psi^{1/\gamma}} (Y^h + S^g), S^g \right\} \tag{17}$$

$$C(S^g, \Gamma) = Y^h + S^g - S(S^g, \Gamma). \tag{18}$$

This is shown visually in the left panel of Figure 3. For high values of S^g , services S are identical to S^g and, as shown in the right panel of Figure 3, the felicity of the government and households are the same. However, once S^g falls below the fiscal emergency threshold of $\psi^{-1/\gamma} Y^h$, services S are strictly greater than S^g because services are partly financed by a reduction in private household consumption. Because the government does not value this loss to household consumption, government felicity is strictly higher than household felicity when S^g is below the fiscal emergency threshold.

25. There may be multiple values of B' that satisfy the budget constraint due to the endogenous bond price. In this event, the lowest such B' is always optimal as this weakly increases future cash on hand.




FIGURE 3

Comparing government and household felicity

Notes: This figure shows public services and government and household felicity as functions of government spending S^g for a fixed Γ . The left panel shows services $S(S^g, \Gamma)$ with the dashed black line showing the 45-degree line. The right panel shows government felicity $u(S(S^g, \Gamma))$ (dark blue) and household felicity $u(S(S^g, \Gamma)) + v(C(S^g, \Gamma))$ (light yellow). In the right panel, the dashed black line shows the threshold for a fiscal emergency, which is the value of S^g below which transfers from households are positive and $S(S^g, \Gamma)$ no longer falls on the 45-degree line.

How does this affect government decisions? When choosing its portfolio (A', B') ,²⁶ the government is choosing its distribution for future X' . This means it is also choosing the distribution for future $S^{g'}$, i.e. $S^g(X', B', \Gamma', y')$. Relative to households, the government is less averse to distributions that include low $S^{g'}$. This means that, relative to households, the government has less incentive to choose portfolios that avoid states with low $S^{g'}$, such as portfolios with higher total savings $A' - q(B', \Gamma, y)B'$ and portfolios that save in safe bonds rather than the risky asset.

Importantly, this difference in felicity between the government and households becomes more relevant for government decisions as the likelihood of a future fiscal emergency increases and the possibility of receiving transfers becomes more salient. This is a form of gambling for resurrection, where governments in poor fiscal health have an incentive to take additional risk because they are partly insured against outcomes with low $S^{g'}$ but fully enjoy the benefits of outcomes with high $S^{g'}$, as shown in the left panel.

To give a concrete example, a major part of the Detroit, Stockton, and San Bernardino bankruptcies is that these cities issued large amounts of bonds and invested the proceeds into their pension funds. These cities, which were already struggling, bet that their risky pension funds would outperform bonds and became bankrupt after the 2008 stock market bust. Over 70% of all bonds eliminated in U.S. city bankruptcies since 2008 were pension obligation bonds (POB's), i.e. bonds issued explicitly to invest in pension funds. Additionally, Munnell *et al.* (2010) and Munnell *et al.* (2014) show that governments in financial stress and with high debt are the most likely to issue POB's. This is also consistent with the empirical findings of risk shifting in pension funds by Pennacchi and Rastad (2011), Mohan and Zhang (2014), Bradley *et al.* (2016), and Andonov *et al.* (2017).

26. Note that, given (A', B') , the values for all of the government's other choice variables are pinned down by the constraints of equation (16).

In Section 3, I quantify the model parameters that control the difference in felicity between the government and households in fiscal emergencies to match the panel of California city finances. This incentive to gamble for resurrection will be important for explaining why large cities, who have higher debt and are more likely to declare a fiscal emergency, increase their risk exposure more than small cities.

3. QUANTIFICATION

The key aspect of the model is that when the government is close to a fiscal emergency, it has an incentive to decrease savings and increase risk exposure. To quantify the magnitude of this undersaving and excess risk-taking, I study the response of cities to an unexpected increase in debt during 2008–10. As shown in Figure 1, pension fund assets drop after the 2008 financial crisis while pension liabilities continue to grow, meaning that city debt increases substantially. Figure 2 shows the financial crisis did not affect all cities equally. As detailed in Section 1.2.3, cities with large populations tend to have bigger pension systems relative to revenue due to the fact that they spend more on services that require pensioned labour. Thus, large cities tend to be more affected by the increase in debt. Based on the response of large and small cities to this shock, I infer the key model parameters that control undersaving and excess risk. Cities are split into two groups based on populations, with 237 cities and 238 cities in the small and large groups, respectively.

In this section, I first discuss the functional forms used for preferences and default costs, and then estimate the exogenous revenue and return processes. I find that idiosyncratic shocks have higher volatility in small cities, but otherwise the exogenous processes for small and large cities are quite similar. I then estimate the 8 model parameters to match 33 moments from the empirical response of cities to the financial crisis and give a detailed analysis of city savings, exposure to risk, and fiscal emergencies. Lastly, I discuss how the model matches both the low number of observed defaults and the high spread on bonds without requiring high creditor risk aversion.

3.1. Functional forms

Creditors have CRRA preferences, like households and the government, with risk aversion γ_m and discount factor β_m . The creditors' discount factor is set to match a nominal risk-free rate of $R^f = 1.04$ and their risk aversion is set so that creditors price the risky asset.²⁷

In the model, the cost of services is normalized so that one unit of services costs 1 unit in current cash and η units in pension benefits. Because the cash cost of services is normalized rather than the full cash and pension cost, this normalization will make services cheaper in small cities, where η is lower. The full cost of a unit of services is $1 + \frac{\lambda}{r^f + \lambda} \eta$, where the second term represents the value of the risk-free pension benefits being paid to employees with $r^f \equiv R^f - 1$. To account for this normalization, I also normalize the value of ψ in both groups of cities,

$$\psi_i = \bar{\psi} \left(1 + \frac{\lambda}{r^f + \lambda} \eta_i \right)^{\gamma-1} \quad (19)$$

where $i \in \{\text{small, large}\}$ and $\bar{\psi}$ will be an estimated parameter.

27. The nominal risk-free rate is chosen to match the standard values of a 2% real rate and 2% inflation.

TABLE 1
Revenue process parameters

g^{exp}	g^{rec}	ρ	σ_{large}	σ_{small}
1.048 (0.002)	0.944 (0.004)	0.375 (0.012)	0.112 (0.004)	0.177 (0.006)

Notes: This table shows the estimated parameters of the city revenue process. The first three columns give the aggregate growth in expansions and recessions, as well as the persistence of the idiosyncratic component of revenue. These are calculated from a non-linear regression of equation (20) on a panel of revenue for 2003–15. The final two columns give the standard deviation of the idiosyncratic shock for large and small cities, which are calculated from the residuals of the regression. The low value for the persistence and the high values for the standard deviations σ_{large} , σ_{small} mean that cities in the data experience large but short-lived deviations from the aggregate trend.

Lastly, the aggregate growth shocks take three values, $g_t \in \{g^{exp}, g^{rec}, g^{dis}\}$, representing an expansion, recession, or rare disaster with probabilities $1 - \pi^{rec} - \pi^{dis}$, π^{rec} , π^{dis} . The possibility of a rare disaster is useful for matching the measured bond spreads in the data and this rare disaster is not realized during the 2008–15 panel. I set the default cost so that $\phi(g^{exp}) = \phi(g^{rec})$, meaning that the default cost will only differ in the disaster state. This is to reflect the fact that in the data, there was not a large increase in defaults during the 2008–10 recession. Most of the defaults occurred after the recession had ended and there has similarly not been a large increase in defaults in previous recessions. The only period where municipal defaults dramatically increased was the Great Depression, which I am treating as a disaster state. During this period, Congress passed laws making municipal bankruptcy and default substantially easier and agents may expect that default costs would again be lowered if another disaster occurs.²⁸

3.2. Exogenous processes

To estimate the revenue process for small and large cities, I use a longer panel of revenue that covers 453 cities from 2003 to 2015. I set 2009 and 2010 as the recession years since average revenue for both groups fell in these years and set all other years as expansion years. The values for growth in expansions and recessions g^{exp} , g^{rec} and the persistence of the idiosyncratic shocks ρ are estimated from the following pooled regression

$$\log(Y_{i,t}) = \alpha_i + \rho \log(Y_{i,t-1}) + (1 - \rho) \log(g^{exp}) t + (\log(g^{rec}) - \log(g^{exp})) [t_{t,2009} + (1 - \rho)t_{t,2010} - \rho t_{t,2011}] + \varepsilon_{i,t} \quad (20)$$

where α_i is a city fixed effect and $t_{t,2009}$, $t_{t,2010}$, $t_{t,2011}$ are dummies indicating if t is greater than or equal to 2009, 2010, and 2011 respectively. The regression can be estimated separately for large and small cities, but the estimated values of g^{exp} , g^{rec} , ρ only slightly differ between the groups. Appendix C.5 provides the derivation of equation (20) from equations (1)–(3) and gives more detail on the separate regressions. Where large and small cities do differ is in the volatility of their idiosyncratic shocks, σ , which is estimated from the residuals of the regression. Table 1 gives the estimated parameters for the revenue process.

The remaining parameters for exogenous processes are the returns, the disaster parameters, the probabilities of the aggregate states, and the maturity probability for pension benefits. Since almost all plans are run through CalPERS, I estimate the returns in expansions and recessions

28. Prior to the Great Depression, cities could not legally declare bankruptcy to seek protection from creditors. In 1934, Congress created the Chapter 9 bankruptcy code for municipalities, which was subsequently followed by the largest series of municipal defaults in U.S. history (Moringiello, 2014).

TABLE 2
Estimated parameters

β	γ	$\bar{\psi}$	y^h	$\phi(\sim g^{dis})$	$\phi(g^{dis})$	η_{large}	η_{small}
0.95	6.61	16.62	1.46	2.17	0.16	1.00	0.67

Notes: This table shows the estimated parameters that best fit 33 targeted moments. These moments are the average debt-to-revenue, pension assets-to-revenue, and pension liabilities-to-revenue for large and small cities in each year from 2011 to 2015, as well as the number of cities with fiscal emergencies, the number of defaulting cities, and the average bond spread. Since the model is overidentified, I use a diagonal weighting matrix where the elements are the inverse of the standard deviation of the estimated moments. The first two parameters are the discount factor and risk aversion for the government. The third and fourth parameters control the household preference for consumption and the level of their disposable income. The fifth and sixth parameters are the default cost as a fraction of revenue outside of the disaster state and in the disaster state. The seventh and eighth parameters are the accrual rates for benefits for large and small cities.

from the CalPERS return for 2003 to 2015, which gives $R(g^{exp}) = 1.100$ and $R(g^{rec}) = 0.928$. The parameters for the disaster state are taken from Barro (2006). To show that the model can generate large spreads and support a high equity premium without needing large disasters, I use conservative estimates for g^{dis} , $R(g^{dis})$, π^{dis} of 0.85, 0.85, 0.015. The probability of recession π^{rec} is estimated from the NBER recession series using all post-1945 observations. Controlling for the fact that the 1.5% chance of disaster is not realized over this period, the probability of recession is 0.143. Finally, λ is set to 0.025 to match the empirical duration of pension benefits of 13 years. Appendix D derives the formula for duration in the model.²⁹

3.3. Parameters and moments

The eight model parameters are the preference parameters β , γ , $\bar{\psi}$, the level of household disposable income that households can give to the government y^h , the default costs $\phi(\sim g^{dis})$, $\phi(g^{dis})$, and the accrual rate for pension benefits for large and small cities η_{large} , η_{small} . Table 2 shows the estimated values. These parameters are estimated to match six time series, the number of defaulting cities, the number of cities with fiscal emergencies, and the average bond spread for 2010–15. The six time series are the average debt, risky assets, and pension liabilities, all relative to revenue, for large and small cities. Combined, a city's debt, pension assets, and pension liabilities define its entire balance sheet. So, matching these six time series means that the model will accurately capture how each aspect of the balance sheet changes over time for both large and small cities.

Given a set of parameters, I use the following method to estimate the model response to the financial crisis. First, I solve the model described in Section 2 for both large and small cities to get the policy functions for the government, the bond price function, and the transfer function. Second, I use the estimated shock process and the realized revenue to infer the realized exogenous shocks g_t , Γ_t , and $y_{i,t}$. Last, I set each city's initial state variables to their 2010 measured values and then feed in the realized exogenous shocks to calculate the model government's choice of next period pension assets, pension liabilities, and bonds, as well as the bond price (if any bonds are issued) and whether a fiscal emergency is declared.

In the data, the pension liabilities are measured as the present value of benefits discounted at $r^p = 7.5\%$ and the debt is the value of unfunded pensions and bonds outstanding. Thus, the

29. As discussed in Section 1.5, the empirical duration of 13 years is measured from the sensitivity of CalPERS pension liabilities to changes in the discount rate.

TABLE 3
Estimated moments

	Data	Model
Defaulting cities	0.42%	0.42%
Fiscal emergency cities	12.00%	12.00%
Average bond spread	1.08%	1.08%

Notes: This table shows the values for three targeted moments in the data and in the model. The first two rows show the percent of cities that default and the percent of cities that declare a fiscal emergency during 2010–15. The third row shows the average bond spread. In the data, the tax-adjusted spread is measured as $yield_{i,t}/(1 - \tau_t) - r_t^f$, where τ_t is the marginal tax rate and r_t^f is the risk-free rate at time t measured from daily Treasury yield curves. Schwert (2017) estimates that 74% of the tax-adjusted spread on municipal bonds is due to default premia. In the model, bond spreads are entirely due to default premia, so I use 74% of the average tax-adjusted spread in the data as my target for model bond spreads. Data on bond spreads comes from MSRB municipal bond trades for California city bonds.

model equivalent of these variables is

$$L = \frac{\lambda}{r^p + \lambda} B^p, \quad (21)$$

$$D = B + L - A. \quad (22)$$

The percent of defaulting cities for 2010–2015 in the data is 0.42% (San Bernardino and Stockton), the percent of cities with fiscal emergencies is 12% (57 cities), and the targeted spread on bonds is 1.08%. Table 3 and Figure 4 summarize the fit of the estimated model. While the model and the data have the same average pension fund return in expansions, the data realizations for returns fluctuate each year while the model returns for 2011–15 are always $R(g^{exp})$. This is why the paths for debt and pension assets are more volatile in the data than in the model.

Appendix C provides a more detailed discussion of the parameter identification, including the Jacobian of moments with respect to each of the eight parameters and a model version of Figure 2. To give the general intuition, the discount factor β and the risk aversion γ are primarily determined by the paths for debt and risky assets, respectively, across both groups of cities. The values for $\bar{\psi}$, y^h are controlled by the difference between small and large cities in debt and risky assets, as well as the number of fiscal emergencies, since these parameters control the transfer function. As we will see in Section 3.5, large cities are more likely to declare a fiscal emergency over this period and therefore their choices for debt and risky assets are more sensitive to changes in the transfer function than the choices of small cities. As discussed in Appendix C, the choice of large city debt will depend more on y^h , while $\bar{\psi}$ will mainly impact the choice of large city risky assets. The values for η_{large} , η_{small} are determined by the paths of liabilities for the two groups and the default costs $\phi(\sim g^{dis})$, $\phi(g^{dis})$ primarily depend on the number of defaults and the average bond spread.

3.4. Response of savings and risk exposure

Given that undersaving and excess risk are the main inefficiencies in the model, the key variables of interest are the debt and risk exposure of cities. To make comparison across groups and time easier, Figure 5 shows the change in debt-to-revenue and pension assets-to-revenue since 2010. As seen in both Figures 4 and 5, cities have gradually decreased debt since 2010, but even after five years of aggregate expansion shocks, cities have only undone about 1/4 of the increase in debt from 2008 to 2010. For small and large cities, debt-to-revenue increased 45pp and 63pp respectively from 2008 to 2010 and has only declined 9pp and 17pp from 2010 to 2015. Over




FIGURE 4
Paths of debt, pension assets, and pension liabilities

Notes: This figure shows the average value of the three components of balance sheets for small and large cities in each year. The dashed black line separates the 2008–10 shock from the 2010–15 response to the shock. The left panel shows the average debt-to-revenue for small cities (light yellow) and large cities (dark blue). The solid lines show the values from the data for 2008–15. The dotted lines show the model values for average debt-to-revenue for small (light yellow) and large (dark blue) cities for 2010–15. The middle panel shows the average pension assets-to-revenue for small and large cities in the data and in the model using the same format. The right panel shows the average pension liabilities-to-revenue for small and large cities in the data and the model, again using the same format.

this same time, exposure to risk has increased for both groups. Pension assets-to-revenue have increased 26pp and 34pp for small and large cities. The fact that risky assets have increased more than debt has decreased means that cities also have bigger balance sheets, *i.e.* the increase in assets is offset by an increase in bond and pension liabilities.

Comparing across groups, large cities have a bigger increase in debt from 2008 to 2010. This is because they tend to run larger pension systems and were more affected when pension assets declined while pension liabilities continued to grow. In response, large cities decrease debt 8pp more than small cities from 2010 to 2015, but they do not do this by issuing fewer bonds or accumulating fewer pension liabilities. As shown in the right-hand side of Figure 5, this is entirely done by increasing pension fund assets 8pp more than small cities.³⁰ Because we have hundreds of cities in each group, these 8pp differences are both statistically significant. Thus, across time and across groups, cities are slowly reversing the debt increase from the financial crisis, primarily by investing more in their risky pension funds.

3.5. Response of fiscal emergencies

After the 2008–10 shock, the number of cities declaring fiscal emergencies increases dramatically. For 2000 to 2007, only 14 cities declare a fiscal emergency. For 2010–15, more than four times as many cities, 57, declare a fiscal emergency. What is not obvious beforehand is which

30. This is not explained by large cities simply earning more on their pension fund assets. In both the model and the data, the pension fund earns higher than average returns for 2010–15 due to consistent aggregate expansion shocks. However, this only accounts for 2.7pp of the 8pp difference. The primary source of the difference is that large cities are putting more money into their funds.




FIGURE 5
Change in debt and pension assets

Notes: This figure shows the average change in debt-to-revenue and pension assets-to-revenue from 2010. The dashed black lines separate the 2008–10 shock from the 2010–15 response to the shock. The left panel shows the average change in debt-to-revenue from 2010 for small cities (light yellow) and large cities (dark blue) for 2008–15. The solid lines show the values from the data for 2008–15. The dotted lines show the model values for the average change in debt-to-revenue for small (light yellow) and large (dark blue) cities for 2010–15. The right panel shows the average change in pension assets-to-revenue from 2010 for small and large cities in the data and in the model using the same format.

types of cities will declare fiscal emergencies. Large cities run bigger pension systems and were more exposed to pension fund losses during the crisis. However, large cities also have more stable revenues ($\sigma_{large} = 0.112$, $\sigma_{small} = 0.177$), meaning they may be able to spread these losses over time to avoid any major service reductions. In comparison, small cities may be forced into a fiscal emergency due to large negative idiosyncratic shocks.

As shown in Table 4, the model correctly predicts that large cities will be more prone to fiscal emergencies over this period. In the data, 14.29% of large cities declare a fiscal emergency while only 9.70% of small cities declare a fiscal emergency. In the model, 16.39% of large cities declare fiscal emergencies, compared to only 7.59% of small cities, meaning that exposure to pension fund losses dominates the differences in revenue volatility. Combined with the results in Figure 5, this means that large cities are more likely to declare a fiscal emergency but are still choosing to increase their risk exposure more than small cities.

3.6. Defaults and bond spreads

As documented in the literature on municipal bonds, the frequency of default for municipal bonds is quite low outside of the Great Depression, yet these bonds still carry large default risk premia. I calculate the bond spread using the yield for all trades on non-insured, tax-exempt bonds for the 475 cities over 2010–15. Specifically, I measure the tax-adjusted spread as $yield_{i,t}/(1 - \tau_t) - r_t^f$, where τ_t is the marginal tax rate at time t and r_t^f is the risk-free rate at time t measured from the daily Treasury yield curve. I calculate the marginal tax rate for creditors using the top California state income tax bracket and the top federal income tax bracket in each year. Since my model uses one-year bonds, I use all trades for bonds with maturity of 9–15 months and estimate an average tax-adjusted spread of 1.46% (0.06%).

TABLE 4
Fiscal emergencies by city type

	Large	Small	All
Data	14.29%	9.70%	12.00%
Model	16.39%	7.59%	12.00%

Notes: The first and second columns show the percent of large and small cities that declare fiscal emergencies during 2010–15 in the data and in the model. These moments are not targeted when estimating the model parameters. The model correctly predicts that large cities will have more fiscal emergencies over this period than small cities. The third column shows the percent of all cities declaring fiscal emergencies, which is targeted when estimating model parameters.

As argued in [Stock \(1994\)](#), [Liu et al. \(2003\)](#), and [Schwert \(2017\)](#), this spread can be primarily explained by default risk premia. [Schwert \(2017\)](#) estimates that 74% of the tax-adjusted spread is due to default risk premia. Since spreads in my model are solely due to default risk, I target a model spread of 1.08% to match 74% of the total 1.46% spread.

To match the low number of defaults after a recession and the high default risk premia, the model gives a high value for the default cost in expansions and recessions and a low value in disasters. Because creditors price the risky asset, the high spread cannot be due to creditors having extremely high risk aversion. In other words, the risk aversion required to generate this spread with few defaults in all states would not be consistent with the equity premium, since the pension fund is mainly equity.³¹ Therefore, the spread must be due to a small chance that many cities will simultaneously default if a disaster is realized, similar to the large series of defaults seen after the Great Depression. This effectively makes the municipal bonds state-contingent contracts, where the cities will almost always pay the full amount B in non-disaster states but may pay a lower amount $\phi(g^{dis})Y$ in a disaster.

4. RESULTS

Using the quantified model, I study the risk of future fiscal emergencies and how government policies can be changed to increase household welfare. This is done in two steps. First, in Sections 4.1 and 4.2, I estimate how many cities would declare fiscal emergencies if another stock market bust occurred and find that many cities are unprepared for this negative shock. I then calculate the spending and risky assets that a benevolent government would choose and how average city debt-to-revenue and pension assets-to-revenue would have changed from 2010 to 2015 under these choices.

Second, motivated by the fact that a benevolent government chooses lower spending and saves more in bonds, in Sections 4.3 and 4.4, I analyse two policy experiments: (i) savings requirements and (ii) requiring the pension fund to invest more in risk-free bonds. Savings requirements can produce large welfare gains for households while a more bond-heavy pension fund actually produces welfare losses because it reduces the government's incentive to save. The

31. For robustness I test the effects of using higher creditor risk aversion and introducing bankruptcy inefficiency where creditors do not receive the full default cost that is paid by the city. In equilibrium, high creditor risk aversion tends to lead to lower average spreads, as cities choose to issue fewer or no bonds. Because cities are relatively patient ($\beta R^f = 0.99$), they will not borrow if the high spreads are mainly due to disagreement between creditors and cities about the risk-neutral probability of future states. Similarly, large bankruptcy inefficiency also leads cities simply to not issue bonds and contradicts the high recovery rate for municipal bonds seen in the data.

savings requirement is conceptually and quantitatively distinct from a pension funding requirement, as a pension funding requirement does not limit bond issuance and produces much smaller welfare gains for households.

Thinking about the practical implementation of these policy experiments, two details are important: these requirements would likely have to be imposed on cities by the county or state government as the city governments would not voluntarily choose to restrict themselves in these ways, and these policy experiments assume that retroactive pension benefit increases are also banned.³² Taking the benefits as given, the goal is to understand whether household welfare can be increased while still fulfilling the promised pension obligations. Section 4.5 provides a deeper discussion of implementation.

4.1. *Stock market bust*

Since 2009, pension funds have consistently delivered positive annual returns, even for 2020 and 2021, and stocks in general have experienced a persistent bull market. However, it is unlikely that this will continue forever, which begs the question of how well cities can handle another negative shock to pension fund assets. To answer this, I estimate the effect of a -24% shock to pension fund assets, consistent with the net effect of the 2008–10 shock. I initialize each city to its 2015 empirical state variables, reduce pension fund assets by 24% , and then simulate 10,000 paths for the future exogenous shocks g_t , $y_{i,t}$ that cities receive. For each path, I calculate the number of cities that declare a fiscal emergency and then average across paths.

Given the recent financial crisis, one might expect that cities would alter their savings and risk exposure to better prepare for negative shocks in the future. However, I find that the hypothetical 2015 shock causes more than twice as many cities to declare a fiscal emergency as the original 2008–10 shock. As shown in Figure 6, on average 136 cities declare a fiscal emergency in the first six years after the 2015 shock, which is over 25% of all cities. For comparison, in both the model and the data, 57 cities declare a fiscal emergency in the first six years after the 2008–2010 shock. Just as in the empirical and model response to the 2008–10 shock, most of the cities declaring fiscal emergencies after the hypothetical 2015 shock are large cities, 81 out of 136. Because of this, these 136 cities with fiscal emergencies cover 48% of California's population. In other words, nearly half of all people in California live in cities that cannot handle another stock market bust without declaring a fiscal emergency.

Looking at Figures 4 and 5, we can see the reason why. For both large and small cities, debt is substantially higher in 2015 than in 2008. As discussed in Section 3.4, cities have only undone about $1/4$ of the 2008–10 increase in debt-to-revenue. This means that the 2015 shock hits cities that are starting much closer to the fiscal emergency region than the 2008 cities. On top of this, pension assets relative to revenue are higher in 2015 than in 2008 for both groups, meaning that the same percent loss on the pension fund will have a larger impact on cities' balance sheets. In short, cities are more indebted and more exposed to risk in 2015 than 2008.

While the data on individual city pension systems is limited to 2015, I can show that continued positive returns on the pension fund for 2016–22 will not reverse this result. I consider a simple scenario where no new pension benefits are accrued and no money is added to the pension fund after 2015. The pension fund simply earns the observed CalPERS return r_t for 2016–22 and maturing benefits are paid with money from the fund. Let $\tilde{B}_{i,t}^p$, $\tilde{A}_{i,t}$, $\tilde{L}_{i,t}$ denote

32. Given the current, heavily underfunded status of these pensions, a retroactive benefit increase seems unlikely. However, for the long-term implications of the savings requirement, it is important that higher savings is not simply offset by higher retroactive benefits.




FIGURE 6

Cities with fiscal emergencies within 6 years of stock market bust

Notes: This figure shows the number of cities that declare a fiscal emergency in the first six years after a shock. Large cities are shown in dark blue and small cities are shown in light yellow. The numbers above the bars indicate the number of large cities declaring a fiscal emergency and the total number of cities declaring a fiscal emergency. The first column shows the number of cities declaring fiscal emergencies in the data for 2010–15 after the 2008–10 shock. The second column shows the number of cities declaring a fiscal emergency in the model for 2010–15 after the 2008–10 shock. The third column shows the number of cities declaring a fiscal emergency for 2015 to 2020 after a hypothetical -24% shock to pension assets in 2015, matching the net effect of the 2008–10 shock.

the pension benefits, pension fund assets, and pension liabilities in this scenario. The benefits and assets follow $\tilde{B}_{i,t+1}^p = (1 - \lambda)\tilde{B}_{i,t}^p$ and $\tilde{A}_{i,t+1} = (1 + r_t)\tilde{A}_{i,t} - \lambda\tilde{B}_{i,t}^p$. Using equation (21) to substitute benefits for liabilities, we can derive the following,

$$\tilde{L}_{i,t+1} - \tilde{A}_{i,t+1} = (\tilde{L}_{i,t} - \tilde{A}_{i,t})(1 + r^p) + (r^p - r_t)\tilde{A}_{i,t}. \quad (23)$$

In words, this says that unfunded pensions $\tilde{L}_{i,t} - \tilde{A}_{i,t}$ will grow by $r^p = 7.5\%$ each year and will only decrease if the pension fund return r_t exceeds the discount rate for the liabilities r^p . For 2016–22, the annualized return on the pension fund was 6.71%, so unfunded pensions will actually be higher in 2022 than in 2015, unless cities specifically act to put more money into the funds. So, recent returns on the pension fund are not enough to decrease city debt and move cities away from the fiscal emergency region. Further, this result is not driven by events in 2020–22. The annualized return on the pension fund for 2016–19 was 6.78% which still implies that unfunded pensions would be higher in 2019 than in 2015.

4.2. Benevolent government

Given the high risk of future fiscal emergencies, a natural question is whether this is optimal for households. To test this, I calculate the policy functions for a benevolent government that has the same preferences (6) as households. Because the benevolent government and households have the same preferences, we can collapse the two stages of the subgame in Section 2.3 into a single stage in which households make all of the decisions. Specifically, households choose spending, the future portfolio, and the transfers S^h , subject to the constraint that transfers cannot

be negative,

$$\begin{aligned}
 \hat{V}(X, B^P, \Gamma, y) &= \max_{S^g, S^h, A', B', B^{P'}} u(S^g + S^h) + v(Y^h - S^h) + \beta \mathbb{E} \left[\hat{V}(X', B^{P'}, \Gamma', y') \mid \Gamma, y \right] \\
 S^g &= X - A' + q(B', \Gamma, y) B' \\
 B^{P'} &= (1 - \lambda) B^P + \eta(S^g + S^h) \\
 S^h &\geq 0 \\
 A' &\geq 0 \\
 X' &= \Gamma' y' - \lambda B^{P'} - \min \{B', \phi(g') \Gamma' y'\} + R(g') A'.
 \end{aligned} \tag{24}$$

To see how much these differences in policy functions matter, I calculate the response of debt and risky assets under the benevolent policy functions to the sudden increase in debt in 2008–10 and compare this to the actual response. I then measure the welfare gains for households from switching to the benevolent government's policy functions.

4.2.1. Benevolent versus baseline policy functions. As discussed at the end of Section 2.3, the baseline government's choices can be boiled down to its choice of spending S^g (or equivalently savings $X - S^g$) and its choice of risk exposure A' . Holding S^g fixed, an increase in exposure to stock market risk A' means the government has the same total savings but is choosing higher risk exposure by borrowing more in bonds and saving more in the risky pension fund. Similarly, holding A' fixed, an increase in spending S^g means that the government has the same exposure to stock market risk but is choosing lower total savings by borrowing more in bonds. Figure 7 shows how the baseline government's choice of spending and risk exposure compares to the values chosen by the benevolent government. Both graphs show the policies as functions of cash on hand X for a fixed (B^P, Γ, y) with the dashed line indicating the threshold value of cash on hand where the baseline government declares a fiscal emergency.

The left panel of Figure 7 shows that, compared to the benevolent government, the baseline government spends more when not in a fiscal emergency. When cash on hand is above the FE threshold, the baseline government is not currently in a fiscal emergency, but it knows that it can declare a fiscal emergency next period if future cash on hand falls below the threshold due to negative shocks. Since the baseline government does not value the cost to households of providing fiscal emergency transfers, it has a lower desire to save and spends more than the benevolent government would choose. As cash on hand decreases and approaches the threshold, the gap between baseline and benevolent spending increases, as the likelihood of a fiscal emergency in the future becomes higher. In other words, overspending is worse when the government is near a fiscal emergency. As cash on hand increases, the baseline and benevolent choices for spending converge.

Conversely, when cash on hand is below the threshold, the baseline government is currently receiving transfers from households and knows that it may not receive transfers next period if positive shocks push future cash on hand above the threshold. This incentivizes the baseline government to spend less than the benevolent government would choose. In general, overspending is more frequent than underspending since cities spend the majority of their time not in a fiscal emergency.

Overall, the left panel of Figure 7 highlights an important conceptual point. The agency friction (*i.e.* the difference in felicity between the baseline government and households) induces a “feast or famine” behaviour for the baseline government. In good states, the baseline government chooses higher spending than the benevolent government. Because the baseline government does




FIGURE 7

Baseline and benevolent government policy functions

Notes: This figure shows the policy functions of the baseline government (dark blue), which does not care about household consumption, and the benevolent government (light yellow), which does care about household consumption and maximizes household utility. The policies are all shown as functions of cash on hand for a fixed (B^P, T, y) . In the left panel, the solid lines show the choice of spending for both governments. The dashed line indicates the cash on hand threshold for a fiscal emergency for the baseline government. The benevolent government also declares fiscal emergencies, but spending does not discontinuously change around this threshold. In the right panel, the solid lines show the choice of risk exposure A' for both governments and the dashed line again indicates the cash on hand threshold for a fiscal emergency for the baseline government.

not save enough in good states, when a bad shock eventually hits the city, the baseline government substantially cuts spending S^g , declares a fiscal emergency, and asks households for a transfer. Thus, the agency problem is not just about pension fund risk exposure, but also excess spending in good states which keeps cities inefficiently close to the fiscal emergency region and vulnerable to negative shocks.

The right panel of Figure 7 shows the choice of stock market risk exposure A' . Importantly, while pension fund assets A' are a form of savings, the left panel of Figure 7 already tells us how baseline government savings compare to benevolent government savings. This is because savings $X - S^g$ are just the complement of spending S^g . Thus, the right panel of Figure 7 does not tell us about total savings but instead tells us whether the government is primarily saving through risky pension fund assets or through bonds.

While the baseline government has an incentive to choose greater risk exposure than the benevolent government, particularly when close to the fiscal emergency threshold, it is limited by the fact that it can only borrow at high spreads. This means that at low cash on hand, when both the baseline and benevolent governments are choosing to borrow with bonds, they choose similar values for risky assets. The two governments still have different portfolio compositions, as the differences in S^g shown in the left panel imply that the baseline and benevolent government choose different B' even when their choices of A' are similar. The largest differences between baseline and benevolent choices for A' occur at high values for cash on hand, where the benevolent government would choose to save in both bonds and risky assets, while the baseline government chooses to continue borrowing in bonds and saving solely in risky assets. This difference in policy functions at high cash on hand is even larger if non-tradable pension benefits B^P are large, since this pushes the government closer to a fiscal emergency while still keeping cash on hand high.




FIGURE 8

Change in debt and pension assets under benevolent policy functions

Notes: This figure shows the average change in debt-to-revenue and pension assets-to-revenue from 2010 under the benevolent government policy functions. The dashed black lines separate the 2008–10 shock from the 2010–15 response to the shock. The left panel shows the average change in debt-to-revenue from 2010 for small (light yellow) and large cities (dark blue). The solid lines show the values from the data for 2008–15. The dotted lines show the model values for the average change in debt-to-revenue for small (light yellow) and large (dark blue) cities for 2010–15 with a benevolent government. The right panel shows the average change in pension assets-to-revenue from 2010 for small and large cities in the data and in the model with a benevolent government using the same format.

To assess these differences in policy functions quantitatively, Figure 8 shows the response of debt and risky assets to the sudden increase in debt in 2008–10 under the benevolent government's policy functions. The solid lines show the actual response of debt-to-revenue and assets-to-revenue in the data and the dashed lines show the counterfactual response if governments were benevolent. The largest difference is that benevolent governments save substantially more. While large and small cities in the data decrease debt-to-revenue by 17pp and 9pp respectively, benevolent governments decrease debt-to-revenue for the two groups by 59pp and 49pp. This is enough to more than offset the 2008–10 increase for small cities and almost completely offset the increase for large cities. The right-hand side of Figure 8 shows that this is not simply due to benevolent governments increasing risky assets substantially more than in the data. Instead, they decrease debt by also issuing fewer bonds.

Comparing across groups, benevolent governments decrease debt-to-revenue more for large cities than for small cities, but not by saving more in risky assets. In the data, large cities decrease debt-to-revenue 8pp more than small cities by increasing assets-to-revenue 8pp more. Benevolent governments instead choose for both groups to increase assets-to-revenue by virtually the same amount and for large cities to lower bonds-to-revenue by 9pp more than small cities. In contrast to raising assets, issuing fewer bonds reduces large cities' levered exposure to risk.

Importantly, benevolent governments initially choose to increase risky assets less for large cities than for small cities. Because large cities are at higher risk of a fiscal emergency and have higher debt, benevolent governments choose for large cities to take less risk and save mainly by issuing fewer bonds. It is not until 2014, when benevolent governments have decreased debt 9pp more for large cities than small cities that benevolent governments choose for large cities to increase risky assets as much as small cities. In other words, benevolent governments closer to a fiscal emergency reduce levered risk by issuing fewer bonds, while governments in the data that are closer to a fiscal emergency increase levered risk by investing more into the risky assets.

TABLE 5
Household welfare gains from changing policy functions

	Benevolent government			Simple cap			Optimal cap		
	Large	Small	All	Large	Small	All	Large	Small	All
% Consumption	0.98	0.87	0.93	0.49	0.64	0.57	0.79	0.78	0.79
% Services	2.41	2.04	2.23	1.25	1.60	1.42	1.96	1.92	1.94

Notes: This table shows the welfare gain for households from changing government policy functions. The welfare gain is measured as the percent increase in consumption (services) in all states and periods that produces the same increase in household utility (6). The welfare gain is calculated for each city starting at its 2010 empirical state variables. The columns show the average for large, small, and all cities. In *Benevolent government*, the baseline government is replaced with a benevolent government that values private consumption. In *Simple cap*, government spending is limited to essential services plus interest on cash on hand. *Optimal cap* uses a stricter cap of essential services plus 0.4% of cash on hand. The optimal cap generates 85% (0.79/0.93) of the benevolent government welfare gain for all cities.

4.2.2. Welfare gains. The first three columns of Table 5 show the welfare gains from switching to the benevolent government, assuming each city starts at its 2010 empirical state variables. This is measured as the percent increase in consumption in all states and periods that produces the same increase in household utility (6). The welfare gain is also measured as a percent of services using the same method. The benevolent government produces sizable welfare gains both as a percent of consumption and as a percent of services, averaging 0.93% and 2.23%, respectively. This is surprising given that changes in the policy functions do not have any first-order effects (e.g. Lucas, 1987). They only alter the choice about spending now versus later, the choice of risk exposure, and the composition of services versus private consumption.

The reason the welfare gains are not trivial is because the policy functions of the baseline government ensure that cities that start near the fiscal emergency region will consistently remain near this region. As shown in Figure 7, cities near the fiscal emergency threshold underspend when in an emergency and overspend when not in an emergency. As discussed above, this “feast-famine” pattern produces excess volatility in services compared to what a benevolent government would choose. It also means that cities currently near the fiscal emergency threshold will continue to be near the threshold in the future, as they save less when cash on hand is above the threshold and save more when current cash on hand is below the threshold. Thus, these cities will continue to have excessively volatile spending that is too high after good shocks and too low after bad shocks.

An important caveat is that these welfare gains are based on the estimated risk aversion of $\gamma = 6.61$. These second-order effects would be smaller under a lower risk aversion coefficient. Rephrased, in an environment with risky assets and realistic risk premia, we need high risk aversion in order to explain portfolio choices, which implies that agents care about second-order effects such as the “feast-famine” pattern in services.

As a concrete example, Figure 8 shows that after the 2008–10 shock to debt, cities only slightly decreased debt-to-revenue. Because of this, many cities are still quite close to the fiscal emergency region and will cut services enough that households will provide transfers if another negative shock occurs, as discussed in Section 4.1. In comparison, benevolent governments would choose to decrease debt-to-revenue considerably more and would be better prepared for a future negative shock. In other words, cities are spending more during the positive expansion shocks of 2011–15 and will have to spend less if negative shocks are realized, while benevolent governments would choose spending that is less volatile across shock realizations.

Comparing across groups, the welfare gains are higher for large cities. In large cities, the cost of services is more backloaded, since services require a higher amount of pension benefits.

This means that saving is more important for large cities and any issues with overspending will be amplified. Similarly, because large cities have more safe pension liabilities, it is easier for them to have levered risk exposure. While benevolent governments would choose to offset these additional safe pension liabilities by issuing fewer bonds, the cities can choose to hold more risky assets instead. In comparison, small cities can only achieve the same levered risk exposure by issuing more bonds to increase their risky assets, which is discouraged by the high spread on bonds.

4.3. *Savings requirements*

Given that benevolent governments would choose to lower debt-to-revenue more than the cities in the data, a natural policy experiment is a state law imposing savings requirements. For the model, a savings requirement is equivalent to a spending cap, since any cash on hand that is not spent must be saved in the pension fund or saved/borrowed in bonds. I use the level of essential services as part of the spending cap. In practice, this is the threshold for services that determines if a city can legally declare a fiscal emergency. In the model, this will be $\theta \equiv \psi^{-1/\gamma} Y^h$, where cities declare a fiscal emergency if and only if spending falls below this threshold. I consider spending caps of the form

$$S^g \leq \kappa X + \theta. \quad (25)$$

Table 5 shows the results for two spending caps, assuming all cities start at their 2010 empirical state variables. The first is a simple cap with $\kappa = E[R^f/g] - 1 = 1.2\%$. This says that city spending is limited to the level of essential services plus the interest on cash on hand. The interest is normalized to account for the fact that there is growth in the model. Even this simple policy produces large welfare gains that are more than half of the total welfare gain of switching to the benevolent government.

To test how close spending caps can get to the benevolent government welfare gain, I calculate the optimal κ to maximize the welfare gain for all cities. This produces a stricter cap of $\kappa = 0.4\%$, which limits spending to essential services plus a smaller portion of cash on hand than the simple cap. Since overspending is a bigger problem for large cities, the stricter cap primarily raises the welfare gain for large cities, from 0.49% of consumption to 0.79%, but small cities still see a noticeable improvement from 0.64% of consumption to 0.78%. Averaging across cities, this optimal cap produces welfare gains of 0.79% of consumption and 1.94% of services. This means that households would be willing for the government to throw away almost 2% of its revenue each year if it also implemented this spending cap. Discounting those future revenues at $R^f = 1.04$, this would be over \$139 billion.

The optimal spending cap generates 85% of the welfare gain of switching to a benevolent government. This means we can get quite close to the optimum just by forcing cities to save more. As discussed above, the baseline government's incentives become more skewed as it gets closer to the fiscal emergency region. By forcing cities to save more, spending caps move them away from the fiscal emergency region. As the likelihood of a fiscal emergency in the future becomes smaller, the baseline government's incentives align more closely to those of the benevolent government and any inefficiencies become smaller. Quantitatively, these benefits outweigh the fact that spending will be low as cities transition from their 2010 state variables to lower levels of debt and higher cash on hand, and that the spending cap will reduce flexibility in the government's choice of spending.

Conveniently, we can also extend these results to a model in which households have a higher β (i.e. greater patience) than the baseline government. Changes to the household β do not change the transfer function of equation (12), which means it does not change the baseline government's

problem (16). Thus, changes to household β would only alter the benevolent government's problem (24). Mechanically, increasing the household β worsens the problem of undersaving (*i.e.* the baseline government saving too little compared to the benevolent government) which increases the relevance of imposing savings requirements.

4.3.1. Savings requirement versus pension funding requirement. How does the savings requirement compare to a pension funding requirement? Conceptually, the key difference is that a savings requirement also incorporates bonds. A pension funding requirement forces cities to invest more in their pension funds, however, cities may choose to issue bonds in order to increase their pension fund assets. This means cities would not actually be saving more and, as discussed in Section 2.5, high debt governments already engage in this type of behaviour. In comparison, a savings requirement forces cities to save more in their pension fund or save more (borrow less) in bonds. Issuing bonds to invest in the pension fund will not help cities meet the savings requirement.

Conveniently, we can not only compare the savings requirement to a pension funding requirement, but also to a continuum of policies that lie between a savings requirement and a pension funding requirement. Given the state variables (X, B^p, Γ, y) , a pension funding requirement,

$$A' \geq G(X, B^p, \Gamma, y) \quad (26)$$

for some function $G(X, B^p, \Gamma, y)$, is equivalent to a restriction on spending of

$$S^g \leq H(X, B^p, \Gamma, y) + q(B', \Gamma, y) B' \quad (27)$$

where $H(X, B^p, \Gamma, y) \equiv X - G(X, B^p, \Gamma, y)$. This is almost identical to a savings requirement (*i.e.* a spending cap), except for the fact that the limit on spending depends on the state variables and the choice of bond issuance B' . This means that we can write a general spending constraint

$$S^g \leq H(X, B^p, \Gamma, y) + \alpha_{fund} q(B', \Gamma, y) B' \quad (28)$$

where $\alpha_{fund} = 0$ corresponds to a savings requirement and $\alpha_{fund} = 1$ corresponds to a pension funding requirement. Framing the constraint in this way means that we can also explore intermediate policies $\alpha_{fund} \in (0, 1)$.

To test if the distinction between a savings requirement and a pension funding requirement is important, I take the spending caps $H(X, B^p, \Gamma, y) = \kappa X + \theta$ studied in Table 5 and estimate the welfare gain for households for values of α_{fund} ranging from 0 to 1. As shown in Figure 9, changing completely from a savings requirement ($\alpha_{fund} = 0$) to a pension funding requirement ($\alpha_{fund} = 1$) entirely eliminates the welfare gain. More importantly, even small values of α_{fund} , such as 0.1, substantially reduce the welfare gain.

Figure 9 highlights that an effective policy should target a city's combined pension and bond finances. Even a small ability to skirt the savings requirement by issuing more bonds gives the city government enough freedom to undo most of the welfare gains. In a similar sense, Appendix E shows that shifting from a savings requirement to a bond issuance cap also eliminates the welfare gain, as it simply encourages the city government to spend more by saving less in their pension funds. Intuitively, because money is fungible, policies that only target the choice of bonds or the choice of pension assets are less effective than a policy which targets the joint bond and pension asset choice.




FIGURE 9

Welfare gain of savings requirement versus pension funding requirement

Notes: This figure shows the change in welfare for households for various values of α_{fund} . Following equation (28), values of α_{fund} from 0 to 1 cover a continuum of policies from a pure savings requirement ($\alpha_{fund} = 0$) to a pure pension funding requirement ($\alpha_{fund} = 1$). The dark blue line shows the welfare gain starting from the simple spending cap at $\alpha_{fund} = 0$ and transitioning to a pension funding requirement. The light red line shows the welfare gain starting from the optimal spending cap at $\alpha_{fund} = 0$ and transitioning to a pension funding requirement. The welfare gain is measured for each city starting at its 2010 empirical state variables and then averaged across cities.

4.4. Safer pension fund

Since the two inefficiencies are undersaving and excess risk-taking, the second logical policy experiment is to reduce the risk of the pension fund assets. In Section 4.1, we see that many cities would declare a fiscal emergency if the pension fund experienced a large negative shock. In this section, I test whether household welfare can be improved by forcing the pension fund to invest more in risk-free bonds to prevent these types of large losses. Specifically, I change the pension fund return from its current value of $R(g)$ to

$$\tilde{R}(g) = \alpha_{rf} R^f + (1 - \alpha_{rf}) R(g), \quad (29)$$

where α_{rf} is the weight on risk-free bonds. Since cities can only issue defaultable bonds that have large spreads, they cannot simply undo this change by issuing their own risk-free bonds and putting more money into the fund.

Figure 10 shows the change in household welfare in percentage points of consumption for all cities for different levels of α_{rf} . There are no levels of α_{rf} that produce non-trivial welfare gains and many values produce sizable welfare losses. Starting with the positive values of α_{rf} , we see that making the pension fund safer actually reduces welfare. Specifically, household welfare monotonically decreases as α_{rf} increases, declining by 0.01% of consumption at $\alpha_{rf} = 0.1$ and by 0.39% of consumption at the completely risk-free pension fund where $\alpha_{rf} = 1$. This is because a safer pension fund discourages cities from saving due to the lower average return. In other words, the benefits of reducing excess risk-taking are quantitatively dominated by the fact that this worsens undersaving.

Given that increasing α_{rf} reduces welfare by worsening the problem of undersaving, we may actually want to decrease α_{rf} to induce cities to save more. In the data, the pension fund invests roughly 20% into bonds, so one can think of $\alpha_{rf} = -0.2$ as a policy where the pension fund sells all of its bonds and invests purely in the risky portions of the fund. Values of α_{rf} below -0.2




FIGURE 10

Welfare gain from changing pension fund risk

Notes: This figure shows the change in welfare for households from investing the pension fund into safer assets. The line shows the welfare change in percentage points of consumption from changing the pension fund return from its current value of $R(g)$ to a weighted sum of the current return and the risk-free rate $\alpha_{rf}R^f + (1 - \alpha_{rf})R(g)$. The welfare change is measured for each city starting at its 2010 empirical state variables and then averaged across cities.

would correspond to the pension fund buying call options on stocks or taking other measures to increase exposure to excess returns. The left portion of Figure 10 shows that as we decrease α_{rf} below 0, we do get welfare increases and the slope with respect to α_{rf} is initially quite similar to the slope for positive α_{rf} . However, as we continue decreasing α_{rf} , the welfare gain levels off and begins falling as excess risk-taking becomes a larger problem. This means that the welfare gains are limited to small values of less than 0.01% of consumption.

4.5. Implementation of a savings requirement

Given the effectiveness of the savings requirement relative to a restriction on risk, a relevant question is whether a savings requirement can be feasibly implemented. In practice, households may not be able to commit to enforcing a restriction during crises. The model and empirical evidence on fiscal emergencies already highlight a key example of this limitation. California state laws impose strong restrictions on city taxes, yet during crises city governments are allowed to raise temporary emergency taxes and households cannot commit ex ante to deny these emergency funds. Would the savings requirement of equation (25) suffer from a similar lack of enforceability?

Fortunately, the savings requirement is not intended to bind in the worst states. Recall that a savings requirement is equivalent to a spending cap and, as shown in the left panel of Figure 7, the baseline government does not choose excessive spending in the lowest cash on hand states. The baseline government only spends more than the benevolent government when cash on hand is above the fiscal emergency threshold. The purpose of the spending cap is not to reduce spending in all states, but rather to remove the “feast-famine” pattern in the baseline government’s spending wherein spending is too high in good states and too low in bad states. Thus, the spending cap is designed to be slack in bad states and only bind in good states. Mechanically, this

comes from the constant in the spending cap of equation (25), which allows the government to spend $\theta > 0$ even if cash on hand is zero.³³

Because the savings requirement is not designed to bind during the worst states, the welfare gains of Table 5 do not require households to be able to enforce the spending cap during crises. In fact, the estimated welfare gains already incorporate that, ex post, households may approve additional services. The spending cap is imposed on S^g rather than total services S . Even in states where the constraint on S^g binds, households will always approve emergency tax increases S^h to fund additional services if the marginal utility of services exceeds the marginal utility of consumption. The purpose of the policy experiment is to demonstrate that even with these limitations on ex ante commitment, the savings requirement is still effective in achieving the majority of the welfare gains of completely switching to a benevolent government.

Additionally, this aspect of being lax in the worst states highlights the key distinction between the spending cap and another policy which could potentially increase savings, namely a balanced budget requirement. While there is some degree of creative accounting, Bohn and Inman (1996) and Poterba (1996) both find that state-level balanced budget requirements significantly impact governments' tax and spending choices.³⁴ A common tradeoff studied in this literature is that a balanced budget requirement makes government spending procyclical by limiting spending in bad states (e.g. Clemens and Miran, 2012). In contrast, the spending cap is designed to make spending less procyclical (i.e. to reduce or eliminate the "feast-famine" pattern in spending), making it potentially more palatable to households as it does not require cutting essential services during crises.

In terms of legislative implementation, the existing restrictions on taxes provide a potential blueprint for the spending cap (i.e. the savings requirement). Current laws such as Proposition 13 and Proposition 218 demonstrate that (i) households can propose and pass laws through the ballot initiative that significantly restrict the actions of local governments and (ii) these restrictions carry significant bite in practice (e.g. Shires, 1999; Rueben, 2003). Just as those laws limit city governments' ability to approve new taxes, a new ballot initiative could limit the ability to approve new hiring or new projects whenever spending exceeds the prescribed cap. Like the tax restrictions, the spending cap could include an exemption for fiscal emergencies, wherein households are allowed to approve emergency tax increases to fund additional government spending, consistent with the treatment of fiscal emergencies in the model policy experiment.

5. ALLOWING PRIVATE SAVING AND BORROWING

The results of Section 4 show that undersaving, or equivalently overspending, is the main source of inefficiency. Figure 8 shows that a benevolent government would save substantially more than the baseline government, and Table 5 shows that imposing a spending cap can move the economy almost completely to the benevolent outcome. The standard Ricardian equivalence prediction would be that any overspending by the government can be completely negated by private savings, making households indifferent to the choice of spending. To test this prediction,

33. For intuition, in the case of $\kappa = 0$, the savings requirement simply becomes an upper bound on possible spending, $S^g \leq \theta$, which makes it straightforward that the constraint is slack in bad states and only binding in good states.

34. Quoting Poterba (1996), "While some cosmetic changes are used to meet state balanced-budget requirements, these changes are quantitatively less important than tax increases and spending cuts." Bohn and Inman (1996) similarly find that end-of-year balanced budget requirements for the general fund have a positive effect on savings and do not find that the requirements lead to governments shifting their deficits into other non-general-fund accounts.




FIGURE 11

Change in debt and pension assets

Notes: This figure shows the average change in debt-to-revenue and pension assets-to-revenue from 2010. The vertical lines separate the 2008–10 shock from the 2010–15 response to the shock. The left panel shows the average change in debt-to-revenue from 2010 for small cities (light yellow) and large cities (dark blue) for 2008–15. The solid lines show the values from the data for 2008–15. The dotted lines show the values from the model without private savings. The dashed lines show the values from the model with private savings. The right panel shows the average change in pension assets-to-revenue from 2010 for small and large cities in the data, in the model without private savings and in the extended model with private savings using the same format.

I study an extended version of the model where households can privately save and borrow. The full description of this model is provided in Appendix A.

In this extended model, Ricardian equivalence does not hold, meaning that the government's choices do impact household welfare and cannot be negated by private saving or borrowing. This occurs because there are two goods, private consumption and public services, and households cannot trade claims on public services. When the government overspends and current services are too high, households would like to sell current services but they have no way to do this. Increasing private savings in response to government overspending will just lower current consumption rather than lower current public services. In short, private saving and borrowing does not allow households to undo government overspending.

Additionally, I find that the quantitative predictions of the model with private savings are almost the same as those from the model without private savings. In other words, allowing households to borrow and save does not significantly change the government's behaviour as households have little ability to undo the government's actions. Thus, for the purposes of studying government behaviour, having households that simply consume or transfer their income is a reasonable approximation of the outcome when households privately save/borrow and significantly simplifies a complicated dynamic game between households and the government.

As in Section 3.4, I use the policy functions to calculate the response of cities to the 2008–10 shock. I initialize household bonds to 0 in 2010. This is done to ensure that differences in government policy between the two models are not due to the model households simply starting the period richer or poorer in the model with private savings than the model without private savings.

Despite the ability of households to privately borrow and save, the level of public services is almost unchanged. Compared to the model without private savings, average public services for 2010–15 are 0.26% higher in the model with private savings. For reference, under a benevolent government average public services would be 7.02% lower. In other words, allowing households to privately borrow and save does not prevent inefficiently high public services. Figure 11 shows

the change in debt-to-revenue and pension assets-to-revenue since 2010 in the data, in the model without private savings, and in the model with private savings. Compared to the model without private savings, governments in the model with private savings decrease debt slightly more slowly and accumulate slightly less in risky assets.

6. CONCLUSION

How do governments manage large pension systems and how does this affect fiscal emergency risk? This paper shows that governments act as levered investors, borrowing in bonds and pension liabilities and investing in risky pension funds. As such, they are highly exposed to stock market risk and may turn to households for transfers after bad shocks. I focus on California cities, where I find that pension liabilities and assets are an order of magnitude larger than bonds outstanding and pension liabilities are non-defaultable. This makes pension liabilities and assets an important part of city finances which should be incorporated when assessing fiscal health.

To understand government pension decisions, I propose a model of governments that not only choose their spending and saving/borrowing in bonds but also their pension liabilities and savings in risky pension funds. The ability of the government to receive transfers in a fiscal emergency gives it an incentive to save less and take more risk, particularly when a fiscal emergency becomes more likely. This can explain the response of California cities to the financial crisis, where cities choose low savings and substantially increase risk exposure. In particular, these skewed incentives are important for matching the finding that large cities, which are more likely to declare a fiscal emergency, increase risk exposure more after the crisis than small cities. Because of this increase in risk exposure and low savings, cities are vulnerable to another stock market bust. I estimate that a hypothetical bust at the end of the sample would lead to more than twice as many fiscal emergencies as the 2008 financial crisis.

Savings requirements can address these incentive issues and achieve almost all of the welfare gains for households from moving to a constrained optimum with a benevolent government. This is because savings requirements reduce the problem of both undersaving and excess risk-taking. As the government accumulates more savings, the likelihood of a future fiscal emergency decreases and the effect of fiscal emergency transfers on government incentives lessens. As a result, savings requirements also reduce the government's incentive to take excess risk. In contrast, forcing the pension fund to invest in safer assets only addresses risk-taking at the expense of savings. A safer pension fund decreases the government's exposure to risk, but the lower average return on the fund reduces the government's incentive to save. In the quantified model, the benefits of reducing excess risk-taking are outweighed by the cost of lowering savings and the policy results in a welfare loss for households.

Acknowledgments. I thank Adrien Auclert, Luigi Bocola, Hanno Lustig, Monika Piazzesi, Joshua Rauh, Martin Schneider, and Christopher Tonetti for their generous advice as well as Jose Barrero, Gideon Bornstein, Gila Bronshtein, Laura Castillo-Martinez, Scarlet Chen, Ricardo Delao, Paul Dolfin, Eran Hoffman, Patrick Kehoe, Nadia Kotova, John Kramer, Ernest Liu, Alessandra Peter, Ciaran Rogers and seminar participants at the Stanford Macro Lunch for helpful comments. Financial support from the Peter G. Peterson Foundation, the Gerald J. Lieberman Fellowship, as well as the B.F. Haley and E.S. Shaw Fellowship for Economics and the Shultz Graduate Student Fellowship in Economic Policy through grants to the Stanford Institute for Economic Policy Research is gratefully acknowledged.

Supplementary Data

Supplementary data are available at [Review of Economic Studies](#) online.

Data Availability

The data and code underlying this research are available on Zenodo at <https://doi.org/10.5281/zenodo.21540536>.

REFERENCES

- AGUIAR, M., CHATTERJEE, S., COLE, H., *et al.* (2016), "Quantitative Models of Sovereign Debt Crises", in *Handbook of Macroeconomics* (Vol. 2) (Amsterdam: Elsevier) 1697–1755.
- AGUIAR, M. and GOPINATH, G. (2006), "Defaultable Debt, Interest Rates and the Current Account", *Journal of International Economics*, **69**, 64–83.
- ALFARO, L. and KANCZUK, F. (2009), "Optimal Reserve Management and Sovereign Debt", *Journal of International Economics*, **77**, 23–36.
- ANDONOV, A., BAUER, R. M. M. J. and CREMERS, K. J. M. (2017), "Pension Fund Asset Allocation and Liability Discount Rates", *The Review of Financial Studies*, **30**, 2555–2595.
- ARELLANO, C. (2008), "Default Risk and Income Fluctuations in Emerging Economies", *American Economic Review*, **98**, 690–712.
- ARELLANO, C., ATKESON, A. and WRIGHT, M. (2016), "External and Public Debt Crises", *NBER Macroeconomics Annual*, **30**, 191–244.
- ARELLANO, C. and RAMANARAYANAN, A. (2012), "Default and the Maturity Structure in Sovereign Bonds", *Journal of Political Economy*, **120**, 187–232.
- BARRO, R. J. (2006), "Rare Disasters and Asset Markets in the Twentieth Century", *Quarterly Journal of Economics*, **121**, 823–866.
- BIANCHI, J., HATCHONDO, J. C. and MARTINEZ, L. (2018), "International Reserves and Rollover Risk", *American Economic Review*, **108**, 2629–2670.
- BOHN, H. and INMAN, R. P. (1996), "Balanced-Budget Rules and Public Deficits: Evidence From The US States," in *Carnegie-Rochester Conference Series on Public Policy* (Vol. 45) (Amsterdam: Elsevier) 13–76.
- BOYER, C. (2020), "Public Pensions and State Government Borrowing Costs" (Working Paper).
- BOZ, E. (2011), "Sovereign Default, Private Sector Creditors, and the IFIs", *Journal of International Economics*, **83**, 70–82.
- BRADLEY, D., PANTZALIS, C. and YUAN, X. (2016), "The Influence of Political Bias in State Pension Funds", *Journal of Financial Economics*, **119**, 69–91.
- BRINKMAN, J., COEN-PIRANI, D. and SIEG, H. (2018), "The Political Economy of Municipal Pension Funding", *American Economic Journal: Macroeconomics*, **10**, 215–246.
- BROWN, J. R. and PENNACCHI, G. G. (2016), "Discounting Pension Liabilities: Funding versus Value", *Journal of Pension Economics & Finance*, **15**, 254–284.
- BROWN, J. R. and WILCOX, D. W. (2009), "Discounting State and Local Pension Liabilities", *American Economic Review*, **99**, 538–542.
- CHATTERJEE, S. and EYIGUNGOR, B. (2012), "Maturity, Indebtedness, and Default Risk", *American Economic Review*, **102**, 2674–2699.
- CLEMENS, J. and MIRAN, S. (2012), "Fiscal Policy Multipliers on Subnational Government Spending", *American Economic Journal: Economic Policy*, **4**, 46–68.
- EATON, J. and GERSOVITZ, M. (1981), "Debt with Potential Repudiation: Theoretical and Empirical Analysis", *The Review of Economic Studies*, **48**, 289–309.
- EPPLÉ, D. and SCHIPPER, K. (1981), "Municipal Pension Funding: A Theory and Some Evidence", *Public Choice*, **37**, 141–178.
- FINK, F. and SCHOLL, A. (2016), "A Quantitative Model of Sovereign Debt, Bailouts and Conditionality", *Journal of International Economics*, **98**, 176–190.
- GLAESER, E. L. and PONZETTO, G. A. M. (2014), "Shrouded Costs of Government: The Political Economy of State and Local Public Pensions", *Journal of Public Economics*, **116**, 89–105.
- GORDON, G. and GUERRON-QUINTANA, P. (2024), "On Regional Borrowing, Default, and Migration", *Journal of International Economics*, **150**, 103916.
- HATCHONDO, J. C. and MARTINEZ, L. (2009), "Long-duration Bonds and Sovereign Defaults", *Journal of International Economics*, **79**, 117–125.
- HATCHONDO, J. C., MARTINEZ, L. and ONDER, Y. K. (2017), "Non-defaultable Debt and Sovereign Risk", *Journal of International Economics*, **105**, 217–229.
- INMAN, R. P. (1982), "Public Employee Pensions and the Local Labor Budget", *Journal of Public Economics*, **19**, 49–71.
- LELAND, H. E. (1994), "Corporate Debt Value, Bond Covenants, and Optimal Capital Structure", *The Journal of Finance*, **49**, 1213–1252.
- LIU, S., WANG, J. and WU, C. (2003), "Effects of Credit Quality on Tax-Exempt and Taxable Yields", *The Journal of Fixed Income*, **13**, 80–99.
- LUCAS, R. (1987), *Models of Business Cycles* (New York: Basil Blackwell, New York).
- LUCAS, D. J. and ZELDES, S. P. (2009), "How Should Public Pension Plans Invest?" *American Economic Review*, **99**, 527–532.
- MODIGLIANI, F. and MILLER, M. H. (1958), "The Cost of Capital, Corporation Finance and the Theory of Investment", *The American Economic Review*, **48**, 261–297.
- MOHAN, N. and ZHANG, T. (2014), "An Analysis of Risk-Taking Behaviour for Public Defined Benefit Pension Plans", *Journal of Banking & Finance*, **40**, 403–419.

- MORINGIELLO, J. M. (2014), "Goals and Governance in Municipal Bankruptcy", *Washington and Lee Law Review*, **71**, 403.
- MUMY, G. E. (1978), "The Economics of Local Government Pensions and Pension Funding", *Journal of Political Economy*, **86**, 517–527.
- MUNNELL, A. H., AUBRY, J.-P., CAFARELLI, M., *et al.* (2014), *An Update on Pension Obligation Bonds* (40) (Chestnut Hill, MA: Center for Retirement Research at Boston College).
- MUNNELL, A. H., CALABRESE, T., MONK, A., *et al.* (2010), "Pension Obligation Bonds: Financial Crisis Exposes Risks", *State and Local Issue in Brief*, **9**.
- NOVY-MARX, R. and RAUH, J. D. (2009), "The Liabilities and Risks of State-Sponsored Pension Plans", *Journal of Economic Perspectives*, **23**, 191–210.
- NOVY-MARX, R. and RAUH, J. (2011), "Public Pension Promises: How big are They and What are They Worth?", *The Journal of Finance*, **66**, 1211–1249.
- PANCRAZI, R., SEOANE, H. D. and VUKOTIC, M. (2020), "Welfare Gains of Bailouts in a Sovereign Default Model", *Journal of Economic Dynamics & Control*, **113**, 103867.
- PENNACCHI, G. and RASTAD, M. (2011), "Portfolio Allocation for Public Pension Funds", *Journal of Pension Economics & Finance*, **10**, 221–245.
- POTERBA, J. M. (1996), "Budget Institutions and Fiscal Policy in the us States", *The American Economic Review*, **86**, 395.
- RUEBEN, K. S. (2003), *Fiscal Effects of Voter Approval Requirements on Local Governments* (Public Policy Instit. of CA).
- SCHWERT, M. (2017), "Municipal Bond Liquidity and Default Risk", *The Journal of Finance*, **72**, 1683–1722.
- SHIRES, M. A. (1999), *Patterns in California Government Revenues Since Proposition 13* (San Francisco, CA: Public Policy Instit. of CA).
- STOCK, D. (1994), "Term Structure Effects on Default Risk Premia and the Relationship of Default-Risky Tax-Exempt Yields to Risk-Free Taxable Yields—a Note", *Journal of Banking & Finance*, **18**, 1185–1203.
- TRZCINKA, C. (1982), "The Pricing of Tax-Exempt Bonds and the Miller Hypothesis", *The Journal of Finance*, **37**, 907–923.
- YAWITZ, J. B., MALONEY, K. J. and EDERINGTON, L. H. (1985), "Taxes, Default Risk, and Yield Spreads", *The Journal of Finance*, **40**, 1127–1140.