

Public Employee Pensions and Municipal Insolvency*

Sean Myers[†]

December 27, 2019

[Latest Version]

Abstract

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 spending cap works better than a restriction on risk-taking.

*I am immensely grateful to Monika Piazzesi, Martin Schneider, Hanno Lustig, and Joshua Rauh for their invaluable guidance and support during my PhD. I also thank Adrien Auclert, Luigi Bocola, and Christopher Tonetti for their generous advice as well as Jose Barrero, Gila Bronshtein, Scarlet Chen, Riccardo De la O, Paul Dolfin, Eran Hoffman, Patrick Kehoe, Nadia Kotova, John Kramer, Alessandra Peter, Ciaran Rogers and seminar participants at the Stanford Macro Lunch for helpful comments. Financial support from 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.

[†]Stanford Department of Economics. Contact: seanmyers@stanford.edu, web.stanford.edu/~smyers2.

1 Introduction

Across the US, state and local governments promise large pension benefits to public employees and invest in risky pension funds. Pension liabilities owed to employees totaled \$5.4 trillion in 2018 and pension fund assets totaled \$3.9 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 declared insolvency to ask local households for transfers to maintain essential public services. This begs the question, how do governments choose their levered risk exposure and should government policy be changed to reduce insolvency risk?

I propose a quantitative model of a government that chooses not only its bond saving or borrowing, but also its pension liabilities and investment in a pension fund that earns a risk premium. The main trade-offs are spending today versus saving for the future and saving more in the risky fund versus saving more or borrowing less in bonds. The government provides public services for local households and if services are low enough that households are willing to provide a transfer to pay for additional services, the government can declare insolvency. Matched to a panel of California cities and a hand-collected record of fiscal emergencies, the model estimates that governments are highly vulnerable to another stock market bust due to low savings after the financial crisis and increased exposure to risk. Savings requirements increase household welfare by 0.49-0.77% of consumption, whereas forcing the pension fund to invest in safer securities actually reduces household welfare by 0.01-0.39% of consumption.

At the heart of the model is an agency conflict between households and the government, where the government does not value the cost to households of providing transfers. Compared to a benevolent government that maximizes discounted household felicity, 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 insolvency and transfers become more likely, which can explain why the California cities closest to insolvency take more risk in the data. Savings requirements address these incentive issues because they move the government away from the insolvency region. As the likelihood of insolvency decreases, government incentives become less skewed, which also reduces excess risk-taking. Requiring the pension fund to save in safer securities lowers the average return on the fund, which reduces the government's incentive to save. Thus, it reduces excess risk-taking but worsens undersaving.

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

In the quantified model, the latter force dominates and the result is a welfare loss for households.

Given the size of pension liabilities, I first assess if they actually need to be paid. Using settlement details for all US 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. Because pension liabilities are non-defaultable, bond spreads are imperfect measures of fiscal health. Governments with low outstanding bonds have little incentive to default, even if they have large pension liabilities.² Instead, I measure fiscal health by collecting a novel record of fiscal emergencies, which are a form of service-level insolvency, across all California cities. I then construct a panel of bond and pension finances for all California cities which shows (i) pension assets and liabilities are an order of magnitude larger than bonds outstanding, (ii) cities' total bond and pension liabilities are significantly higher than their pension assets, and (iii) cities are highly levered with large total liabilities and large risky pension assets. There is also substantial heterogeneity across cities, with some having both larger gaps between assets and liabilities and higher levered risk exposure.

To understand how these three findings affect insolvency risk, the model features governments whose revenues and pension fund returns both depend on the realized aggregate shock. Each period, the government chooses how much to spend on services, invest in the pension fund, and borrow or save in one-period bonds. Pension benefits are non-tradable and accrue based on the amount the government spends on services. While the government can default on both benefits 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 insolvency in order to ask households for a transfer is the result of a game between the government and local households. In equilibrium, the government declares insolvency if services fall below a threshold of essential services. Since the government can receive transfers in bad states, it has an incentive to save little for the future and to choose high risk exposure.

A key element of the model is that households have a higher preference for private consumption than the government, which means that the government does not fully value the cost to households of providing transfers in insolvency. I motivate this preference structure in two ways. First, in California, city governments cannot independently change taxes; they must get household approval through a vote. If governments were benevolent, households would give them full control over taxes, so this indicates that

²The mayor of San Diego stated this exact rationale for not defaulting in 2010, as quoted in Section 2.4.

households believe the government would overtax if allowed. Second, this structure for preferences predicts that cities close to insolvency will gamble for resurrection, meaning that they take on more risk because they do not bear the full cost of insolvency. This prediction matches the cross-sectional behavior of California cities, where the cities with the highest bonds outstanding and unfunded pensions take on more risk, and is consistent with previous literature on risk-shifting in pension funds (Pennacchi and Rastad (2011), Mohan and Zhang (2014), Bradley et al. (2016), Andonov et al. (2017)).³ Additionally, 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, consistent with the gambling for resurrection prediction, these cities were making substantial bets on the stock market prior to their bankruptcies.

Since model undersaving and excess risk become worse as cities approach insolvency, I quantify the model to match the response of cities to an unexpected drop in pension 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 labor-intensive services. Splitting cities into two groups based on population, the model successfully matches the savings, risk exposure, and pension decisions of both groups as well as the total number of cities that declare insolvency, the number of cities that default, and the average bond spread. Additionally, the model correctly predicts that large cities will be more likely to declare insolvency, even though small cities receive more volatile revenue shocks. Moreover, the model also correctly predicts that cities with higher pension fund assets relative to revenue before the crisis will be more likely to declare insolvency even after controlling for whether the city is large or small.

Using the panel of city finances and the quantified model, I perform two counterfactuals and two policy experiments. First, I test the ability of cities to handle another stock market bust similar to 2008. 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 insolvency as the original 2008-2010 shock: 136 out of 475 cities compared to only 57. While the panel ends in 2015, I also show that the positive returns on the pension fund for 2016-2019 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 2008-2010 shock, a benevolent government would choose to save substantially

³Unfunded pensions are the difference between pension liabilities and pension assets.

more and would reduce levered risk exposure, particularly for the cities closest to insolvency. Starting at the 2010 empirical state variables, switching to the constrained optimal choices increases household welfare by 0.88% of consumption.

Given that the government undersaves and takes excessive risk, the policy experiments examine savings requirements and reducing the riskiness of the pension fund. A savings requirement is similar to a funding requirement but also includes bond borrowing or saving. While a city can meet a funding requirement by just 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 0.3% of cash on hand increases household welfare by 0.77% of consumption.⁴ Forcing cities to save more moves them away from the insolvency region, which reduces the effect of insolvency transfers on government incentives. As a result, the savings requirement also reduces the incentive to take excess risk and produces 87% of the welfare gain from moving to the constrained optimum.

In contrast, requiring the pension fund to invest more in safe assets reduces household welfare.⁵ Investing the pension fund entirely in risk-free bonds lowers household welfare by 0.39% of consumption. The lower average return reduces the government's incentive to save, and the benefits of reducing risk exposure are quantitatively outweighed by the costs of worsening undersaving. Additionally, I show 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.

Related literature

This paper is related to several strands of literature. The first is a large literature on sovereign default (Eaton and Gersovitz (1981); Aguiar and Gopinath (2006); Arellano (2008); see Aguiar et al. (2016) for a survey). These papers study the borrowing and savings decisions of a government with limited commitment to repaying bonds and show how the endogenous choice of default risk can explain the joint dynamics of income, borrowing, and spreads as well as the frequency of default.

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 for EU nations. The amount of Eurobonds is exogenously capped, and the government

⁴The constant term is chosen as the model threshold for services below which the government declares insolvency.

⁵As 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.

always issues the maximum amount, except during a transition period after the bonds are introduced. Boz (2011), Fink and Scholl (2016), and Pancrazi et al. (2017) analyze non-defaultable loans from international financial institutions such as the IMF that can only be accessed if the government agrees to “conditionality clauses” that restrict their future spending. These papers highlight how access to non-defaultable bonds can lower spreads during times of crisis by reducing the need of the government to borrow in defaultable bonds. In contrast, this paper studies the choice of a government to issue defaultable and non-defaultable debt without caps or restrictions on spending and shows how non-defaultable debt can increase the risk of insolvency. Myers (2017) similarly looks at the effect of pensions on government borrowing at the national level, but focuses on the role of defaultable pay-as-you-go pensions in increasing the government’s ability to commit to repaying bonds.

Another novel feature of the model is that I allow the government to invest 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.

Two recent papers have extended 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. Arellano et al. (2016) and Gordon and Guerron-Quintana (2019) study US states and cities, respectively. Both studies focus on the government’s choice of bonds, while this 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. Additionally, the inclusion of non-defaultable pension benefits and risky pension funds substantially increases cities’ total liabilities and exposure to risk.

My work also builds on an empirical and theoretical literature studying unfunded public employee pensions. Novy-Marx and Rauh (2009, 2011), Brown and Wilcox (2009), and Brown and Pennacchi (2016) show that state and local governments have large unfunded pensions and that the reported values for pension liabilities understate governments' obligations. Pension benefits are promises for safe payments, so 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, I match the stream of benefit payments in the model to data directly, rather than computing the discounted value. Given the size of current unfunded pensions, Novy-Marx and Rauh (2014) calculate that contributions would need to increase substantially to achieve full funding and Anzia et al. (2019) confirm that government pension contributions are rising across the US. Novy-Marx and Rauh (2012) also find that pension fund losses for states increase the spread on their municipal bonds. My paper adds to this literature by examining whether cities are issuing bonds in order to increase pension fund contributions and by studying the effect of large unfunded pensions on insolvency risk.

Since the model government does not value the cost to households of providing transfers, my analysis is closely tied to empirical work documenting risk-shifting in public pension funds. State and local governments delegate the investment of their pension funds to a board of trustees. Pennacchi and Rastad (2011), Bradley et al. (2016), and Andonov et al. (2017) show that boards with more politicians or public employee representatives invest in riskier assets. Importantly, all three papers, as well as Mohan and Zhang (2014), find that boards gamble for resurrection, meaning that they increase risk exposure when benefits are more underfunded or after negative shocks to pension assets. In comparison, private pension plans invest in safer assets when benefits are more underfunded (Rauh (2008)). I examine this same incentive to gamble for resurrection for the government rather than the board. In addition to any excess risk that the board may take by investing the fund mainly in risky securities, governments can increase their risk exposure by issuing their own bonds and putting more money into the pension fund.⁶

On the theoretical front, previous papers have separately modeled the choice of unfunded pensions and the choice of risk exposure. 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) study the government's decision to invest its pension fund in safe or risky assets in a two-period model with

⁶As discussed above, this was a significant part of recent city bankruptcies, with Detroit, Stockton, and San Bernardino all issuing large amounts of bonds specifically to invest in their risky pension funds.

fixed contributions and benefits. Pennacchi and Rastad (2011) extends this choice of risk exposure to an infinitely lived government with an exogenous process for benefits and no 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 insolvency, 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.

Lastly, many papers have studied spreads on municipal bonds and the role of tax restrictions in state and local government borrowing. English (1996) studies the default cost for state governments and several empirical and theoretical papers have shown that the large tax-adjusted spread on municipal bonds can be explained by default premia, despite the low frequency of default (Trzcinka (1982); Yawitz et al. (1985); Stock (1994); Liu et al. (2003); Schwert (2017)). State and local governments often face restrictions on their ability to raise taxes. Wallis (2000) examines how these limitations affect the level of bonds outstanding, Poterba (1994) and Rodden and Wibbels (2010) study the effect on expenditure smoothing, and Poterba and Rueben (1999, 2001) and Johnson and Kriz (2005) analyze the effect on state bond yields. My model studies how governments facing tax restrictions make both their bond and pension decisions and endogenizes the default decision to match the spreads in the data.

Layout

The remainder of the paper is organized as follows. Section 2 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 3 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 insolvency transfers. Section 4 quantifies the model to match the response of cities to the financial crisis. Section 5 estimates the risk of future insolvencies, calculates the constrained optimum with a benevolent government, and examines two policy experiments. Section 6 concludes.

2 Data and Facts

To understand how cities choose their risk exposure and borrowing, and how this affects the chance of insolvency, I will need data on three things: (i) the size and state-contingency of city pension assets, pension liabilities, and bonds outstanding, (ii) city borrowing costs, and (iii) city insolvencies. This section first gives an overview of the

pension system in California cities. Then I present a panel dataset of pension assets, pension liabilities, and bonds outstanding across all California cities for 2008-2015. On average, pension assets and liabilities both dwarf the size of bonds outstanding in all years, making the pension system the main component of city balance sheets. However, there is considerable heterogeneity across cities in the size of their balance sheets (relative to revenue) and the composition of their pension and bond liabilities.

To investigate the state-contingency of pension liabilities relative to bonds, I use data on settlement details for all US city bankruptcies since 2008 as well as historical data for California cities from 1967 to 2019. Across the US, recent bankruptcies show that pension liabilities are less state-contingent than bonds, i.e. they are harder to default on. For California pension liabilities appear virtually non-defaultable. I then detail a novel dataset of fiscal emergencies that I hand-collect for all California cities in order to test predictions on insolvency risk, as well as data on municipal bond spreads that will be used to estimate borrowing costs and default risk. Lastly, I discuss the use of discount rates when calculating pension liabilities and how to account for this by modeling benefit payments rather than a present discounted value.

2.1 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. Over 95% of all California city 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. 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 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. 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.⁷ 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.

Once pension liabilities and pension assets are measured, it is useful to discuss the difference between the two: unfunded pensions. This difference answers the question of how much additional money the city needs to put into its pension fund for it to cover the accrued benefits on average. Even if no new benefits are accrued and the pension fund grows at its average return, unfunded pensions will increase over time if a city does nothing. Specifically, suppose a city has \$1 in unfunded pensions. That dollar was supposed to be in the pension fund earning interest to pay future pension benefits. Unless the city acts to decrease benefits or increase assets, its next year unfunded pensions will be $\$(1 + r)$ dollars. This is analogous to rolling over bond debt and implies that cities that wait to address their unfunded pensions will have to make larger changes to pension assets or benefits in the future.

2.2 Panel of city finances

I construct a panel of annual data on city finances from 2008-2015 for all California cities. 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.⁸ Over 95% of 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. In instances where this report

⁷The 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.

⁸I am grateful to Joe Nation and Pension Tracker for providing access to the pension data.




Figure 1: Time series of city finances

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.

is missing from the SCO's files, I collect the underlying audit from the city clerk's office.⁹ 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 all loan debt in the measure of bonds outstanding.

2.2.1 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. 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%. Thus, the pension system is a large component of city finances which cannot be ignored in an analysis of fiscal health. The finding also suggests that cities are not simply entities that issue bonds to finance services. Instead, cities are more like hedge funds that borrow to invest in risky assets.

The financial crisis left cities with a large gap between pension liabilities and as-

⁹Three cities (Loyalton, Tulelake, 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 Tulelake 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 2013-15 because of a lawsuit with employees.




Figure 2: Cross-section of 2015 pension assets, pension liabilities, and bonds outstanding. This figure shows the pension assets, pension liabilities, bonds outstanding, and net wealth as a percent of revenue across 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 wealth of the city, which is the pension assets minus the pension liabilities and bonds outstanding. The net wealth indicates if the city is an overall net saver or net borrower.

sets. In 2008, pension liabilities and pension assets averaged 149% and 139% of revenue, respectively. The difference – unfunded pensions – was 10% of revenue, smaller than bonds outstanding, which averaged 19% of revenue. However, because the pension fund is invested in risky securities, the pension assets dropped during the financial crisis and the safe pension liabilities continued to grow, creating a large gap between the two. 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 2008.¹⁰ This persistent gap between pension liabilities and assets means that unfunded pensions from 2009-2015 range from 47-64% of revenue.

2.2.2 Cross-section

There is substantial heterogeneity in unfunded pensions 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

¹⁰This same pattern holds across the US. 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.

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 net wealth, which is the value of pension assets minus total liabilities. While wealth as a percent of revenue is -71% on average, this ranges from small, positive values of 2.5% to immense negative values below -300%. These two results both indicate that there are some cities which are at a much higher risk of insolvency than others. Towards the left-side of the graph, there are many cities that have relatively high net wealth 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 also have many cities with low net wealth that 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 mostly likely to become insolvent. 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 2.3 shows that pension liabilities actually appear to be senior to bonds in recent bankruptcies and effectively non-defaultable in California, meaning that large pension liabilities are more likely to drive a city into insolvency than large bonds outstanding.

2.2.3 Large and small city differences

Combining Figures 1 and 2, we know that net wealth decreased 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 4, it will be useful to separate cities that were more and less affected by the decrease in net wealth. 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 labor 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 labor. As a consequence, large cities have higher pension liabilities relative to revenue than small cities. Large cities can cover these liabilities by either

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% (0.3%) of revenue on wages for pensioned labor, compared to only 20.6% (0.3%) for small cities. 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.

2.3 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 US since 2008. 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. Cuts to bonds are taken from Moody's and cuts to pension liabilities are taken from the bankruptcy exit plans. 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 70% of revenue while pension liabilities averaged 339%. In other words, 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 California city 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.¹¹ 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 the CalPERS 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."

The seniority of pension benefits implies that pension plans affect insolvency risk in two ways. First, cities have large unfunded pensions that will need to be paid. Since the liabilities cannot be easily cut and cities do not have enough money in their pension funds to pay their liabilities, cities will have to divert revenue from services. Second, unfunded pensions will grow (shrink) in recessions (expansions). Their levered investment in risky assets tends to perform poorly in recessions, which is when cities have the lowest revenues to cover their losses.

2.4 Insolvencies and bond spreads

In the presence of non-defaultable pension liabilities, bond spreads are a poor measure of a city's fiscal health. The reason is that a city may struggle to pay its large unfunded pensions but will 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."¹²

The measure of fiscal health that I will study is "service-level insolvency." This measure focuses on situations in which the city is unable to provide essential services for residents. This is a better measure than bond spreads, as both bond liabilities and pension liabilities could potentially force city to cut services. It is also more directly tied to household welfare than default risk, as it relates specifically to the public services provided to residents. To measure these insolvencies, I hand-collect a record of fiscal emergencies in California cities since 2000. 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 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 to avoid cutting essential services, such as water, sewage, or public safety.

¹¹Boyer (2019) 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.

¹²Sanders, Jerry, "Debunking the bankruptcy myth", San Diego-Union Tribune (2010).

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.¹³ 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 states:

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 rest of the resolution explains why the city is service-level insolvent. 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 will 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. I collect data on all municipal bond trades where customers are buying or selling to a dealer from MSRB for 2008-2015. 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.

2.5 Discussion of discount rates

Cities report their liabilities as the present value of benefit payments discounted at the average return on the pension fund. 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

¹³Data was collected with the assistance of the Shultz Graduate Student Fellowship in Economic Policy.

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.

To address this issue, the model 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). The fraction of benefits that mature each period is set to match the duration of pension benefits measured in the data. Since 2013, CalPERS plans have reported their liabilities discounted at a risk-free rate as well as the average return on the fund. The duration is estimated based on how much the present value changes when the discount rate is lowered.

3 Model

To understand how city finances affect insolvency risk, and conversely how the risk of insolvency 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 short-lived households, who have preferences over public services as well as their own consumption. 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.

3.1 Endowments, preferences, and assets

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 ,

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

Along with the aggregate component Γ_t , city i revenue has an idiosyncratic component $y_{i,t}$ which 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

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

where $S_{i,t}$ is public services, $u''(S) < 0$, $\lim_{S \rightarrow 0} u'(S) = \infty$, and $\lim_{S \rightarrow \infty} u'(S) = 0$.

Households

Each city i has a series of households that live for one period before being replaced by new households. This is done to capture the fact that households in the data may move between cities, meaning they are only short-term residents of each city.¹⁴ 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 over public services $S_{i,t}$ and household consumption $C_{i,t}$ given by

$$u(S_{i,t}) + v(C_{i,t}). \quad (6)$$

Households have the same felicity for services as the government $u(S_{i,t})$ and their felicity for consumption is increasing and concave, $v'(C_{i,t}) > 0$ and $v''(C_{i,t}) < 0$.

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

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

¹⁴The model outcomes for all city-level variables are identical if households are long-lived and hand-to-mouth.

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

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

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 explain in Section 2.1. 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 dispersed 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 $B_{i,t}^p$.

3.2 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 2.3 that governments fear a "mass exodus" of employees if they default on pension benefits. Because $u'(0) = \infty$, 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 4.1. 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. 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. This would continue until enough creditors deviate that the bonds are competitively priced. Given that payment $\phi(g_t) Y_{i,t}$ to creditors is the only punishment, default on bonds occurs whenever $B_{i,t} > \phi(g_{i,t}) Y_{i,t}$, i.e. the cost of paying the bonds exceeds the default cost. The competitive bond price is then simply the discounted value of future payments,

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

3.3 Game between government and households

I focus on the Markov Perfect Equilibrium of a sequential stage game between households and the government. This equilibrium concept assumes that agents act sequentially, cannot commit to future decisions, and make their decisions based only on payoff-relevant state variables. From here forward, I drop the i and t subscripts.

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 $(A', B', B^{p'})$ for the next period. 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. Formally, we have

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

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

This is designed to capture the fact 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.

Household decision

In the second stage, households choose a transfer S^h if a fiscal emergency was declared by the city government. Given the aggregate state Γ and the government's spending on services S^g , households choose $S^h \geq 0$ to maximize their utility (6) subject to equations

(5), (10), and (11). This gives that the optimal transfer will balance their marginal felicity of services and their marginal felicity of consumption, subject to the constraint that transfers must be positive. Specifically, the transfer will be

$$S^h(S^g, \Gamma) = \begin{cases} z : u'(S^g + z) = v'(Y^h - z) & \text{if } u'(S^g) > v'(Y^h) \\ 0 & \text{otherwise.} \end{cases} \quad (12)$$

The assumptions on $u(\cdot), v(\cdot)$ guarantee that $S^h(S^g, \Gamma)$ exists, is unique, and is the optimal transfer. They also guarantee that transfers are weakly decreasing in S^g and weakly increasing in Γ .

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

$$X = Y - \lambda B^p - \min\{B, \phi(g) Y\} + R(g) A, \quad (13)$$

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

Given the transfer function of households $S^h(S^g, \Gamma)$, the government's problem is

$$\begin{aligned} V(X, B^p, \Gamma, y) &= \max_{S, S^g, \delta, A', B', B^{p'}} u(S) + \beta \mathbb{E}[V(X', B^{p'}, \Gamma', y') | \Gamma, y] \\ S^g &= X - A' + q(B', \Gamma, y) B' \\ S &= S^g + \delta S^h(S^g, \Gamma) \\ B^{p'} &= (1 - \lambda) B^p + \eta S \\ A' &\geq 0 \end{aligned} \quad (14)$$

where the bond price function $q(B', \Gamma, y)$ comes from equation (9). 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 are equal to the non-matured benefits plus newly accrued benefits from services, and lastly, the government cannot borrow in the pension assets. In later sections, it will be useful to discuss how much the government saves, which is simply the amount of cash on hand that is not spent, $X - S^g$.

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. We can simplify the prob-

lem by noting that the government declares a fiscal emergency, $\delta = 1$, if and only if $S^h(S^g, \Gamma) > 0$. This is shown formally in the Appendix. 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 the transfer function is positive.

Given this decision for declaring a fiscal emergency δ , the government just has to choose spending S^g and its risky pension assets A' . The values for $B', S, B^{p'}$ are then determined by the constraints of equation (14), namely the budget constraint, the fiscal emergency transfers, and the law of motion for pension benefits. 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 is saving less or borrowing more in bonds.

3.4 Equilibrium

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

1. Taking as given the transfer function S^h , the policy functions $\{S, S^g, \delta, A', B', B^{p'}\}$ and the value function V solve the government’s problem (14).
2. The transfer function S^h satisfies (12).

3.5 Transfers skew incentives

The main outcome 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. This leads the government to overspend, or equivalently undersave, and take on excess risk relative to a benevolent government that maximizes the discounted series of household felicities

$$\mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t (u(S_{i,t}) + v(C_{i,t})) \right]. \quad (15)$$

Using the decision rule for declaring a fiscal emergency from Section 3.3, $\delta = 1$ if and only if $S^h(S^g, \Gamma) > 0$, we have that equilibrium services and consumption are

$$S(S^g, \Gamma) = S^g + \max \left\{ 0, S^h(S^g, \Gamma) \right\}, \quad (16)$$

$$C(S^g, \Gamma) = Y^h - \max \left\{ 0, S^h(S^g, \Gamma) \right\}. \quad (17)$$

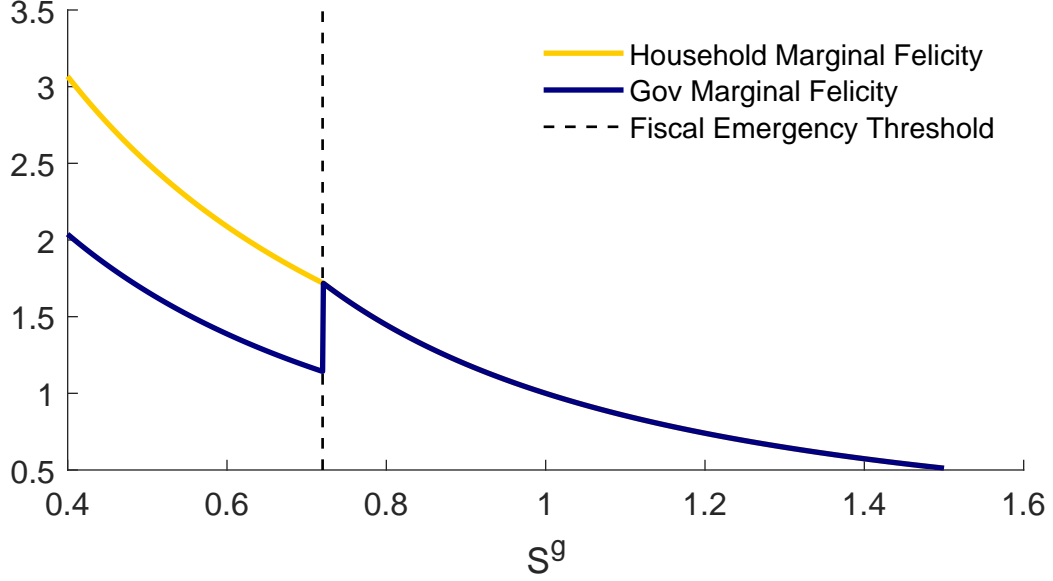


Figure 3: Marginal felicity for government and households

This figure shows the marginal felicity with respect to government spending for households and the government. Given the outcome of the game between households and the government, household marginal felicity (light yellow) is $\frac{\partial}{\partial S^g} [u(S(S^g, \Gamma)) + v(C(S^g, \Gamma))]$ and government marginal felicity (dark blue) is $\frac{\partial}{\partial S^g} u(S(S^g, \Gamma))$. 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. The government's marginal felicity discontinuously jumps when spending falls below the fiscal emergency threshold due to households providing transfers.

Given the outcome of the game between the government and households, Figure 3 shows the marginal felicity for households and the government with respect to S^g . This is done holding Γ fixed. Starting from the right, when spending is high enough that the government does not declare a fiscal emergency, marginal felicity is the same for households and the government and increases as S^g decreases. However, once spending is low enough that the government declares a fiscal emergency, the marginal felicity for the government discontinuously falls. This is because the decrease in services from lowering S^g is partly offset by the increase in transfers from households $S^h(S^g, \Gamma)$. In comparison, the marginal felicity for households continues increasing as S^g falls below the fiscal emergency threshold because these transfers are coming at the expense of consumption. The Appendix shows that for any preferences $u(\cdot), v(\cdot)$ satisfying the assumptions of Section 3.1, the marginal felicity for households strictly increases when spending falls below the threshold for a fiscal emergency, while the marginal felicity for the government decreases if $y^h > 0$, i.e. households have any income to transfer.

Since the government does not value the loss of consumption for households from providing these transfers, the marginal felicity for the government is always strictly below the marginal felicity for households if the government is in a fiscal emergency and

$y^h > 0$. As a result, the government is less concerned than households about bad states where S^g is low, which leads the government to save less for the future and to take on more risk than a benevolent government would choose. Importantly, the government's incentives become more skewed as the likelihood of a future fiscal emergency increases and the possibility of receiving transfers becomes more relevant.

This is a form of gambling for resurrection, where governments close to insolvency have an incentive to take additional risk because they are partly insured against bad outcomes but fully enjoy the benefits of good outcomes. 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 US 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 low net wealth are the most likely to issue POB's.

In Section 4, I quantify the model parameters that control the decline in the government's marginal felicity to match the panel of California city finances. This incentive to gamble for resurrection will be important for explaining why large cities, who have lower net wealth and are more likely to declare a fiscal emergency, increase their risk exposure more than small cities.

4 Quantification

The main outcome 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 after becoming unexpectedly poorer during 2008-2010. As shown in Figure 1, pension fund assets drop after the 2008 financial crisis while pension liabilities continue to grow, meaning that the net wealth of cities declines substantially. Figure 2 shows the financial crisis did not affect all cities equally. As detailed in Section 2.2.3, 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 labor. Thus, large cities tend to be more affected by the increase in unfunded pensions. 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. The model correctly predicts that large cities will be more likely than small cities to declare a fiscal emergency after 2010, even though small cities are more prone to big negative idiosyncratic shocks, and that within each group, cities with higher pension assets relative to revenue in 2008 will be more likely to declare a fiscal emergency. In other words, exposure to the decline in pension fund assets and the increase in pension liabilities significantly raises the probability of a city declaring insolvency after the crisis. 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.

4.1 Functional forms

All agents have CRRA preferences. The government and households have risk aversion γ for services, $S^{1-\gamma}/(1-\gamma)$. Households have the same risk aversion for consumption and a scaling factor ψ , $\psi C^{1-\gamma}/(1-\gamma)$. Creditors have 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.¹⁵

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,

$$\psi_i = \bar{\psi} \left(1 + \frac{\lambda}{r^f + \lambda} \eta_i \right)^{\gamma-1} \quad (18)$$

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

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 $1 - \pi^{rec} - \pi^{dis}, \pi^{rec}, \pi^{dis}$. The possibility of a rare disaster is useful for matching the measured bond spreads in

¹⁵The nominal risk-free rate is chosen to match the standard values of a 2% real interest rate and 2% inflation.

Table 1: Revenue process parameters

g^{exp}	g^{rec}	ρ	σ_{large}	σ_{small}
1.048	0.944	0.375	0.112	0.177
(0.002)	(0.004)	(0.012)	(0.004)	(0.007)

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 (19) on a panel of revenue for 2003-2015. 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.

the data and this rare disaster is not realized during the 2008-2015 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 2008-2010 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.¹⁶

4.2 Exogenous processes

To estimate the revenue process for small and large cities, I use a longer panel of revenue that covers 453 cities from 2003-2015. 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

$$\begin{aligned} \log(Y_{i,t}) = & \alpha_i + \rho \log(Y_{i,t-1}) + (1 - \rho) \log(g^{exp}) t \\ & + (\log(g^{rec}) - \log(g^{exp})) [l_{2009} + (1 - \rho) l_{2010} - \rho l_{2011}] + \varepsilon_{i,t} \end{aligned} \quad (19)$$

where α_i is a city fixed effect and l_{2009} , l_{2010} , l_{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. The Appendix gives more detail on the separate regressions.

¹⁶Prior 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 US history.

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	16.72	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 net average wealth-to-revenue, pension assets-to-revenue, and pension liabilities-to-revenue for large and small cities in each year from 2011-2015, 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.

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 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 $g^{dis}, R(g^{dis}), \pi^{dis}$ 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 a duration of 13 years when moving from the risk-free rate to the discount rate used by city actuaries of 7.5%.¹⁷

4.3 Parameters and moments

The eight model parameters are the preference parameters $\beta, \gamma, \bar{\psi}$, 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 $\eta_{large}, \eta_{small}$. 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

¹⁷Specifically, I set λ so that $\frac{r^p - \lambda}{r^f - \lambda} = \left(\frac{1+r^p}{1+r^f}\right)^{13}$ where the LHS is the ratio of the discounted value at the risk-free rate to the discounted value at the rate used by actuaries $r^p = 0.075$.

Table 3: Estimated moments

	Data	Model
Defaulting cities	2	2
Fiscal emergency cities	57	57
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 number of cities that default during 2010-2015 and the number of cities that declare a fiscal emergency. 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.

emergencies, and the average bond spread for 2010-2015. The six time series are the average net wealth, risky assets, and pension liabilities, all relative to revenue, for large and small cities. Combined, a city's net wealth, 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 3 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 net wealth is the value of pension assets minus pension liabilities and bonds outstanding. Thus, the model equivalent of these variables is

$$L = \frac{\lambda}{r^p + \lambda} B^p, \quad (20)$$

$$W = A - L - B. \quad (21)$$

The number of defaulting cities for 2010-2015 in the data is 2 (San Bernardino and Stockton), the number of cities with fiscal emergencies is 57, and the targeted spread on bonds is 1.08%. Table 3 and Figure 4 summarize the fit of the estimated model. While the model




Figure 4: Paths of net wealth, pension assets, and pension liabilities

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 2008-2010 shock from the 2010-2015 response to the shock. The left panel shows the average net wealth-to-revenue for small cities (light yellow) and large cities (dark blue). The solid lines show the values from the data for 2008-2015. The dotted lines show the model values for average net wealth-to-revenue for small (light yellow) and large (dark blue) cities for 2010-2015. 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.

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 2011-2015 are always $R(g^{exp})$. This is why the paths for wealth and pension assets are more volatile in the data than in the model.

While all 8 parameters are estimated jointly, the discount factor β and the risk aversion γ are primarily determined by the paths for wealth and risky assets, respectively, across both groups of cities. A higher discount factor will cause cities to save more and a higher risk aversion will lower their desire to hold risky assets. The values for $\bar{\psi}$, y^h are controlled by the difference between small and large cities in wealth and risky assets, as well as the number of fiscal emergencies, since these parameters control the transfer function. As we will see in Section 4.5, large cities are more likely to declare a fiscal emergency over this period and therefore their choices for wealth and risky assets are more sensitive to changes in the transfer function than the choices of small cities. A lower value of y^h or a higher value of $\bar{\psi}$ will cause large cities to increase wealth and hold less risky assets, relative to small cities, since they cannot rely on households to provide large transfers if there are bad future shocks. The Appendix details how the

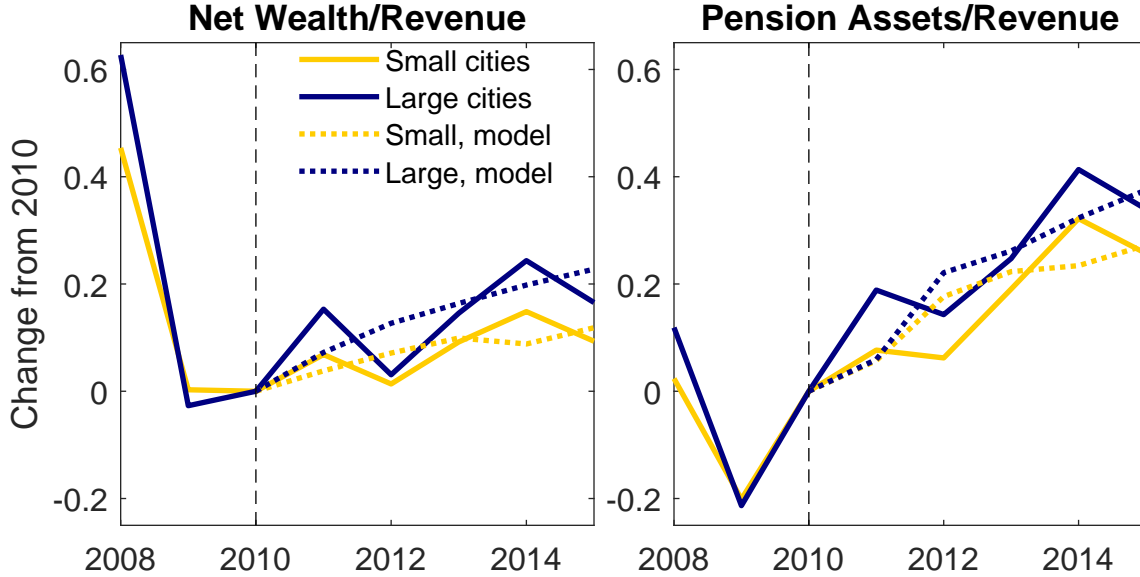


Figure 5: Change in net wealth and pension assets

This figure shows the average change in net wealth-to-revenue and pension assets-to-revenue from 2010. The dashed black lines separate the 2008-2010 shock from the 2010-2015 response to the shock. The left panel shows the average change in net wealth-to-revenue from 2010 for small cities (light yellow) and large cities (dark blue) for 2008-2015. The solid lines show the values from the data for 2008-2015. The dotted lines show the model values for the average change in net wealth-to-revenue for small (light yellow) and large (dark blue) cities for 2010-2015. 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.

choice of wealth will be primarily dictated by y^h while the choice of risky assets will mainly depend on $\bar{\psi}$. 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.

4.4 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 net wealth and risk exposure of cities. To make comparison across groups and time easier, Figure 5 shows the change in net wealth-to-revenue and pension assets-to-revenue since 2010. As seen in both Figures 4 and 5, cities have gradually increased net wealth since 2010. Even after five years of aggregate expansion shocks, cities have only undone about 1/4 of the loss in net wealth from 2008-2010. For small and large cities, net wealth-to-revenue declined 45pp and 63pp respectively from 2008 to 2010 and has only increased 9pp and 17pp from 2010 to 2015. Over 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 in-

Table 4: Fiscal emergencies by city type

	Large	Small	Total
Data	34	23	57
Model	39	18	57

Notes: The first and second columns show the number of large and small cities that declare fiscal emergencies during 2010-2015 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 total number of cities declaring fiscal emergencies, which is targeted when estimating model parameters.

creased more than net wealth 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 decline in net wealth from 2008-2010. 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 increase net wealth 8pp more than small cities from 2010-2015, 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.¹⁸ Thus, across time and across groups, cities are slowly reversing the net wealth decline from the financial crisis, primarily by investing more in their risky pension funds.

4.5 Response of fiscal emergencies

After the 2008-2010 shock, the number of cities declaring fiscal emergencies increases dramatically. For 2000-2007, only 14 cities declare a fiscal emergency. For 2010-2015, more than four times as many cities, 57, declare a fiscal emergency. What is not obvious beforehand is which 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

¹⁸This 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 2010-2015 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.

Table 5: Regression of fiscal emergencies on risk exposure

	Data	Model
2008 $\frac{Assets}{Rev}$	0.045 (0.017)	0.154 (0.019)
1 { <i>Large city</i> }	0.025 (0.030)	0.014 (0.024)

Notes: This table shows that higher values of pension assets-to-revenue before the financial crisis, i.e. higher exposure to risk, are associated with a higher probability of declaring a fiscal emergency after the crisis. The values give the coefficients for a simple regression across cities where the dependent variable is a dummy indicating if the city declared a fiscal emergency during 2010-2015 and the independent variables are the city's 2008 pension assets-to-revenue and a dummy indicating if this is a large or small city. The two columns show that the model correctly predicts that cities with higher assets-to-revenue in 2008 have a higher probability of declaring a fiscal emergency after the 2008-2010 shock, even after controlling for city group.

prone to fiscal emergencies over this period. In the model, 39 of the 57 cities with fiscal emergencies are large, compared to 34 in the data, 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.

We also see that within each group, cities with larger pension fund assets relative to revenue in 2008 were more likely to declare a fiscal emergency. Table 5 shows the results of a simple regression of whether a city declared a fiscal emergency in 2010-2015 on the city's 2008 pension assets-to-revenue, controlling for city group. In both the model and the data, an increase in 2008 pension assets-to-revenue significantly increases the probability of a city declaring a fiscal emergency. In the model, the decision to declare a fiscal emergency only depends on the initial state variables and the series of shocks, whereas in the data, the decision to declare a fiscal emergency may be more noisy. As such, it is not surprising that the coefficient on 2008 assets-to-revenue is higher in the model than the data. What is important is that both coefficients are positive and significant. In other words, we can directly observe that cities with higher exposure to pension fund losses were more likely to become insolvent.

4.6 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 2010-2015. 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%).

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 $0.74 * 1.46\% = 1.08\%$.

To match the low number of defaults for 2010-2015 and the high default risk premia, the model gives a high value for the default cost outside of disasters 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.¹⁹ 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 $\phi(g^{dis})Y$ if a disaster is realized.

5 Results

Using the quantified model, I study the risk of future insolvencies and how government policies can be changed to increase household welfare. First, I estimate how many cities would declare insolvency 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 net wealth-to-revenue and pension assets-to-revenue would have changed from 2010 to 2015 under these choices. Motivated by the fact that a benevolent government chooses lower spend-

¹⁹For 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 to simply not issue bonds and contradicts the high recovery rate for municipal bonds seen in the data.

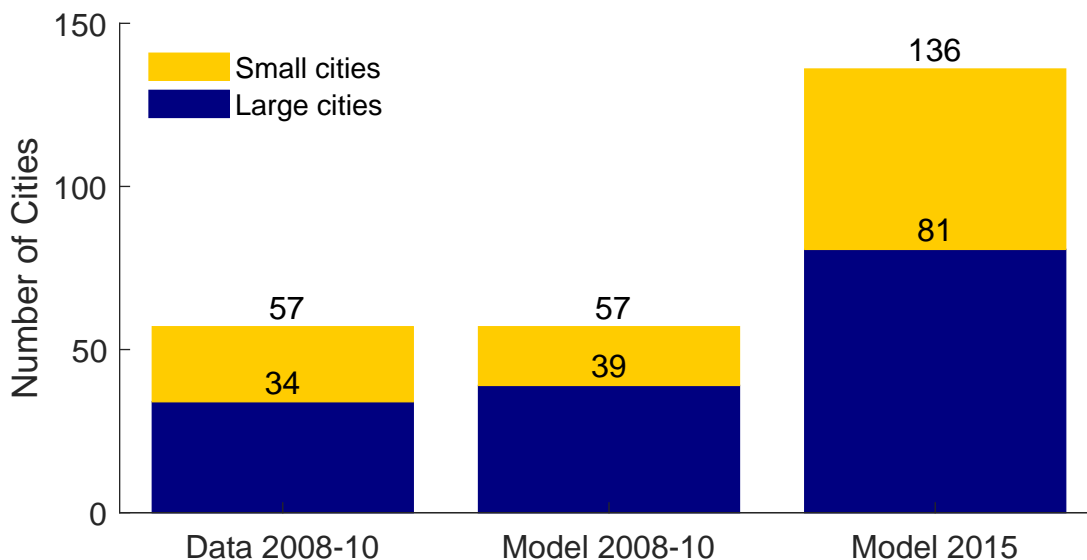


Figure 6: Cities with fiscal emergencies within 6 years of stock market bust

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 2010-2015 after the 2008-2010 shock. The second column shows the number of cities declaring a fiscal emergency in the model for 2010-2015 after the 2008-2010 shock. The third column shows the number of cities declaring a fiscal emergency for 2015-2020 after a hypothetical -24% shock to pension assets in 2015, matching the net effect of the 2008-2010 shock.

ing and saves more in bonds, I analyze 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.

5.1 Stock market bust

Since 2009, pension funds have consistently delivered positive returns, 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 2008-2010 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 2008-2010 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 2008-2010 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, net wealth is substantially lower in 2015 than in 2008. As discussed in Section 4.4, cities have only undone about 1/4 of the 2008-2010 decline in net wealth-to-revenue. This means that the 2015 shock hits cities that are starting much closer to the insolvency 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 poorer 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 2016-2019 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 2016-2019 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 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 (20) to substitute benefits for liabilities, we can derive the following,

$$\tilde{L}_{i,t+1} - \tilde{A}_{i,t+1} = (\tilde{L}_{i,t} - \tilde{A}_{i,t}) (1 + r^p) + (r^p - r_t) \tilde{A}_{i,t}. \quad (22)$$

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 2016-2019, the annualized return on the pension fund was 6.7%, so unfunded pensions will actually be higher in 2019 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 increase city wealth and move cities away from the insolvency region.

5.2 Benevolent government

Given the high risk of future insolvencies, a natural question is whether this is optimal for households. To test this, I calculate the policy functions for a benevolent government that maximizes discounted household felicity (15). This represents a constrained optimum, where the benevolent government is still restricted by the fact that they cannot commit to repay bonds and cannot control the transfer function $S^h(\cdot)$. To get a sense of how much these differences in policy functions matter, I calculate the response of wealth and risky assets under the benevolent policy functions to the sudden decrease in wealth in 2008-2010 and compare this to the actual response. I then measure the welfare gains for discounted household felicity from switching to the benevolent government's policy functions.

5.2.1 Policy functions

Figure 7 shows the spending and risky assets chosen by the baseline government studied in the previous sections and the values that a benevolent government would choose. Both graphs show the policies as functions of cash on hand X for a fixed (B^p, Γ, y) with dashed lines indicating the threshold value of cash on hand where the government declares a fiscal emergency. On the left-hand side, we see that the benevolent government's cash on hand threshold for a fiscal emergency is lower than the threshold for the baseline government. Because the benevolent government values the consumption cost to households from providing transfers, it declares fiscal emergencies in fewer states. The light yellow line shows the spending choice of the benevolent government when cash on hand is above and below the fiscal emergency threshold. Due to the inability to commit to repaying bonds, borrowing is limited and spending declines rapidly as cash on hand decreases.

Compared to the benevolent government, the baseline government spends more when not in a fiscal emergency. The dark blue line shows the spending choice of the baseline government and dashed blue line shows the cash on hand threshold for a fiscal emergency. When cash on hand is above this 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

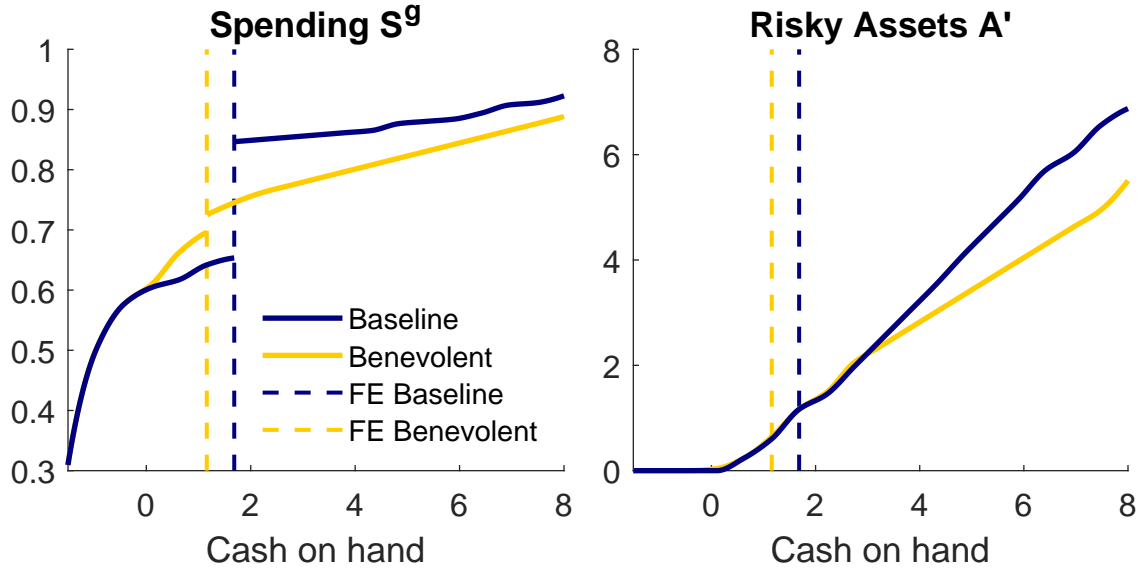


Figure 7: Baseline and benevolent government policy functions

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 discounted household felicity. The policies are all shown as functions of cash on hand for a fixed (B^p, Γ, y) . In the left panel, the solid lines show the choice of spending for both governments and the dashed lines indicate the cash on hand threshold for a fiscal emergency for the baseline and benevolent governments. In the right panel, the solid lines show the choice of pension fund assets A' for both governments and the dashed lines again indicate the cash on hand threshold for a fiscal emergency for both governments.

when the government is near insolvency. 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 will be more important than underspending since cities spend the majority of their time above the threshold, i.e. not in a fiscal emergency.

The right-hand side of Figure 7 shows the choice of risky assets A' . While the baseline government has an incentive to over-invest in risky assets, 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 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




Figure 8: Change in net wealth and pension assets under benevolent policy functions

This figure shows the average change in net wealth-to-revenue and pension assets-to-revenue from 2010 under the benevolent government policy functions. The dashed black lines separate the 2008-2010 shock from the 2010-2015 response to the shock. The left panel shows the average change in net wealth-to-revenue from 2010 for small (light yellow) and large cities (dark blue). The solid lines show the values from the data for 2008-2015. The dotted lines show the model values for the average change in net wealth-to-revenue for small (light yellow) and large (dark blue) cities for 2010-2015 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.

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 insolvency while still keeping cash on hand high.

Figure 8 shows the response of wealth and risky assets to the sudden decrease in wealth in 2008-2010 under the benevolent government's policy functions. The solid lines show the actual response of wealth-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 increase wealth-to-revenue by 17pp and 9pp respectively, benevolent governments increase wealth-to-revenue for the two groups by 59pp and 49pp. This is enough to more than offset the 2008-2010 decline for small cities and almost completely offset the decline 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 increase net wealth by also issuing fewer bonds.

Comparing across groups, benevolent governments increase wealth-to-revenue more for large cities than for small cities, but not by saving more in risky assets. In the data,

large cities increase wealth-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 lower net wealth, 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 increased net wealth 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 insolvency reduce levered risk by issuing fewer bonds, while governments in the data that are closer to insolvency increase levered risk by investing more into the risky assets.

5.2.2 Welfare gains

The first three columns of Table 6 show the welfare gains from switching to this constrained optimum, 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 the discounted sum of household felicity in equation (15). The welfare gain is also measured as a percent of services using the same method. The constrained optimum produces sizable welfare gains both as a percent of consumption and as a percent of services, averaging 0.88% and 2.17%, respectively. This is surprising given that changes in the policy functions do not have any first-order effects. They only alter the choice about spending now versus later and the choice of risk exposure.

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. This 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.

As a concrete example, Figure 8 shows that after the 2008-2010 shock to net wealth,

Table 6: Household welfare gains from changing policy functions

	Constrained optimum			Simple cap			Optimal cap		
	Large	Small	All	Large	Small	All	Large	Small	All
% Consumption	0.93	0.83	0.88	0.49	0.64	0.56	0.77	0.76	0.77
% Services	2.29	2.05	2.17	1.25	1.60	1.42	1.92	1.88	1.90

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 the discounted sum of household felicity (15). 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 *Constrained optimum*, the 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 0.3% of cash on hand. The optimal cap generates almost 90% (0.77/0.88) of the constrained optimum welfare gain for all cities.

cities only slightly increased net wealth-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 5.1. In comparison, benevolent governments would choose to increase net wealth-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 2011-2015 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.

5.3 Savings requirements

Given that benevolent governments would choose to raise wealth-to-revenue more than the cities in the data, a natural policy experiment is a state law imposing savings requirements. This is different from the funding requirements that are typically discussed

when studying pensions because this savings requirement also incorporates bonds. A 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 3.5, governments in financial stress or with low net wealth already engage in this type of behavior. 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. There is also the additional benefit that savings requirements do not force cities to increase their exposure to risk. Rather than requiring that cities invest more in the risky pension fund, cities can meet the savings requirement by either changing their pension fund assets or their bonds.

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

$$S^g \leq \kappa X + \theta.$$

Table 6 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 moving to the constrained optimum.

To test how close spending caps can get to the constrained optimum, I calculate the optimal κ to maximize the welfare gain for all cities. This produces a stricter cap of $\kappa = 0.3\%$, 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 0.77%, but small cities still see a noticeable improvement from 0.64% of consumption to 0.76%. Averaging across cities, this optimal cap produces welfare gains of 0.77% of consumption and 1.90% 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 \$136 billion.




Figure 9: Welfare gain from changing pension fund risk

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 $\alpha R^f + (1 - \alpha) R(g)$. The welfare change is measured for each city starting at its 2010 empirical state variables and then averaged across cities.

The optimal spending cap generates 87% of the welfare gain of switching to a benevolent government. This means we can get quite close to the constrained optimum just by forcing cities to save more. As shown in Section 5.2.1, the government's incentives become more skewed as it gets closer to the fiscal emergency region. This creates a vicious cycle, where overspending is worse when the government is poor, which in turn increases the chance that the government will remain poor in the future. Implementing spending caps reverses this cycle, as cities are forced to save more, which moves them away from the fiscal emergency region. As the likelihood of an 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 higher levels of wealth and cash on hand, and that the spending cap will reduce flexibility in the government's choice of spending.

5.4 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 5.1, we see that many cities would become insolvent 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

$$\tilde{R}(g) = \alpha R^f + (1 - \alpha) R(g),$$

where α 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 9 shows the change in household welfare in percentage points of consumption for all cities for different levels of α . There are no levels of α that produce non-trivial welfare gains and many values produce sizable welfare losses. Starting with the positive values of α , we see that making the pension fund safer actually reduces welfare. Specifically, household welfare monotonically decreases as α increases, declining by 0.01% of consumption at $\alpha = 0.1$ and by 0.39% of consumption at the completely risk-free pension fund where $\alpha = 1$. 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 reduces welfare by worsening the problem of undersaving, we may actually want to decrease α to induce cities to save more. In the data, the pension fund invests roughly 20% into bonds, so one can think of $\alpha = -0.2$ 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 below -0.2 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 9 shows that as we decrease α below 0, we do get welfare increases and the slope with respect to α is initially quite similar to the slope for positive α . However, as we continue decreasing α , 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.

6 Conclusion

How do governments manage large pension systems and how does this affect insolvency 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 insolvency gives it an incentive to save less and take more risk, particularly when insolvency 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 insolvency, 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 insolvencies 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 insolvency decreases and the effect of insolvency 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.

References

- Mark Aguiar and Gita Gopinath. Defaultable debt, interest rates and the current account. *Journal of International Economics*, 69(1):64–83, 2006.
- Mark Aguiar, Satyajit Chatterjee, Harold Cole, and Zachary Stangebye. Quantitative models of sovereign debt crises. In *Handbook of Macroeconomics*, volume 2, pages 1697–1755. Elsevier, 2016.
- Laura Alfaro and Fabio Kanczuk. Optimal reserve management and sovereign debt. *Journal of International Economics*, 77(1):23–36, 2009.
- Aleksandar Andonov, Rob M.M.J. Bauer, and K.J. Martijn Cremers. Pension fund asset allocation and liability discount rates. *The Review of Financial Studies*, 30(8):2555–2595, 2017.
- Sarah F Anzia, Anna Duning Rubens, Lauren Finke, Brad Kent, Tim Tsai, Andrea Lynn, and Devan Shea. Pensions in the trenches: How pension costs are affecting us local government. *Working Paper*, 2019.
- Cristina Arellano. Default risk and income fluctuations in emerging economies. *American Economic Review*, 98(3):690–712, 2008.
- Cristina Arellano and Ananth Ramanarayanan. Default and the maturity structure in sovereign bonds. *Journal of Political Economy*, 120(2):187–232, 2012.
- Cristina Arellano, Andrew Atkeson, and Mark Wright. External and public debt crises. *NBER Macroeconomics Annual*, 30(1):191–244, 2016.
- Robert J. Barro. Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, 121(3):823–866, 2006.
- Javier Bianchi, Juan Carlos Hatchondo, and Leonardo Martinez. International reserves and rollover risk. *American Economic Review*, 108(9):2629–70, 2018.
- Chuck Boyer. Public pensions and state government borrowing costs. *Working Paper*, 2019.
- Emine Boz. Sovereign default, private sector creditors, and the IFIs. *Journal of International Economics*, 83(1):70–82, 2011.
- Daniel Bradley, Christos Pantzalis, and Xiaojing Yuan. The influence of political bias in state pension funds. *Journal of Financial Economics*, 119(1):69–91, 2016.

- Jeffrey Brinkman, Daniele Coen-Pirani, and Holger Sieg. The political economy of municipal pension funding. *American Economic Journal: Macroeconomics*, 10(3):215–46, 2018.
- Jeffrey R. Brown and George G. Pennacchi. Discounting pension liabilities: Funding versus value. *Journal of Pension Economics & Finance*, 15(3):254–284, 2016.
- Jeffrey R. Brown and David W. Wilcox. Discounting state and local pension liabilities. *American Economic Review*, 99(2):538–42, 2009.
- Satyajit Chatterjee and Burcu Eyigungor. Maturity, indebtedness, and default risk. *American Economic Review*, 102(6):2674–99, 2012.
- Jonathan Eaton and Mark Gersovitz. Debt with potential repudiation: Theoretical and empirical analysis. *The Review of Economic Studies*, 48(2):289–309, 1981.
- William B. English. Understanding the costs of sovereign default: American state debts in the 1840's. *The American Economic Review*, pages 259–275, 1996.
- Dennis Epple and Katherine Schipper. Municipal pension funding: A theory and some evidence. *Public Choice*, 37(1):141–178, 1981.
- Fabian Fink and Almuth Scholl. A quantitative model of sovereign debt, bailouts and conditionality. *Journal of International Economics*, 98:176–190, 2016.
- Edward L. Glaeser and Giacomo A.M. Ponzetto. Shrouded costs of government: The political economy of state and local public pensions. *Journal of Public Economics*, 116: 89–105, 2014.
- Grey Gordon and Pablo Guerron-Quintana. On regional borrowing, default, and migration. *Working Paper*, 2019.
- Juan Carlos Hatchondo and Leonardo Martinez. Long-duration bonds and sovereign defaults. *Journal of International Economics*, 79(1):117–125, 2009.
- Juan Carlos Hatchondo, Leonardo Martinez, and Yasin Kursat Onder. Non-defaultable debt and sovereign risk. *Journal of International Economics*, 105:217–229, 2017.
- Robert P. Inman. Public employee pensions and the local labor budget. *Journal of Public Economics*, 19(1):49–71, 1982.
- Craig L. Johnson and Kenneth A. Kriz. Fiscal institutions, credit ratings, and borrowing costs. *Public Budgeting & Finance*, 25(1):84–103, 2005.

- Hayne E. Leland. Corporate debt value, bond covenants, and optimal capital structure. *The Journal of Finance*, 49(4):1213–1252, 1994.
- Sheen Liu, Junbo Wang, and Chunchi Wu. Effects of credit quality on tax-exempt and taxable yields. *The Journal of Fixed Income*, 13(2):80–99, 2003.
- Deborah J. Lucas and Stephen P. Zeldes. How should public pension plans invest? *American Economic Review*, 99(2):527–32, 2009.
- Nancy Mohan and Ting Zhang. An analysis of risk-taking behavior for public defined benefit pension plans. *Journal of Banking & Finance*, 40:403–419, 2014.
- Gene E. Mummy. The economics of local government pensions and pension funding. *Journal of Political Economy*, 86(3):517–527, 1978.
- Alicia H. Munnell, Thad Calabrese, Ashby Monk, and Jean-Pierre Aubry. Pension obligation bonds: financial crisis exposes risks. *State and Local Issue in Brief*, 9, 2010.
- Alicia H. Munnell, Jean-Pierre Aubry, Mark Cafarelli, et al. An update on pension obligation bonds. *Center for Retirement Research at Boston College*, (40), 2014.
- Sean Myers. Sovereign debt, government spending cycles, and back-loaded pension reforms. *Working Paper*, Available at SSRN 2938638, 2017.
- Robert Novy-Marx and Joshua Rauh. Public pension promises: how big are they and what are they worth? *The Journal of Finance*, 66(4):1211–1249, 2011.
- Robert Novy-Marx and Joshua Rauh. The revenue demands of public employee pension promises. *American Economic Journal: Economic Policy*, 6(1):193–229, 2014.
- Robert Novy-Marx and Joshua D. Rauh. The liabilities and risks of state-sponsored pension plans. *Journal of Economic Perspectives*, 23(4):191–210, 2009.
- Robert Novy-Marx and Joshua D. Rauh. Fiscal imbalances and borrowing costs: Evidence from state investment losses. *American Economic Journal: Economic Policy*, 4(2): 182–213, 2012.
- Roberto Pancrazi, Hernán D. Seoane, and Marija Vukotic. Welfare costs of sovereign debt crisis: the role of bailouts. Technical report, 2017.
- George Pennacchi and Mahdi Rastad. Portfolio allocation for public pension funds. *Journal of Pension Economics & Finance*, 10(2):221–245, 2011.

- James M. Poterba. State responses to fiscal crises: The effects of budgetary institutions and politics. *Journal of Political Economy*, 102(4):799–821, 1994.
- James M. Poterba and Kim Rueben. State fiscal institutions and the US municipal bond market. In *Fiscal Institutions and Fiscal Performance*, pages 181–208. University of Chicago Press, 1999.
- James M. Poterba and Kim S. Rueben. Fiscal news, state budget rules, and tax-exempt bond yields. *Journal of Urban Economics*, 50(3):537–562, 2001.
- Joshua D. Rauh. Risk shifting versus risk management: Investment policy in corporate pension plans. *The Review of Financial Studies*, 22(7):2687–2733, 2008.
- Jonathan Rodden and Erik Wibbels. Fiscal decentralization and the business cycle: An empirical study of seven federations. *Economics & Politics*, 22(1):37–67, 2010.
- Michael Schwert. Municipal bond liquidity and default risk. *The Journal of Finance*, 72(4):1683–1722, 2017.
- Duane Stock. 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(6):1185–1203, 1994.
- Charles Trzcinka. The pricing of tax-exempt bonds and the Miller hypothesis. *The Journal of Finance*, 37(4):907–923, 1982.
- John Joseph Wallis. American government finance in the long run: 1790 to 1990. *Journal of Economic Perspectives*, 14(1):61–82, 2000.
- Jess B. Yawitz, Kevin J. Maloney, and Louis H. Ederington. Taxes, default risk, and yield spreads. *The Journal of Finance*, 40(4):1127–1140, 1985.

Appendix

A Model

Two claims are made in Section 3, which both have straightforward proofs.

A.1 Fiscal emergency decision

The first claim is that the optimal choice for declaring a fiscal emergency δ in the government's problem (14) is $\delta = 1$ if and only if the transfer from households is positive, $S^h(S^g, \Gamma) > 0$. If $S^h(S^g, \Gamma) = 0$ then the government is indifferent between $\delta = 0$ and $\delta = 1$. As stated in Section 3.3, I break the indifference by assuming the government does not declare a fiscal emergency when indifferent. Thus, it only needs to be shown that $\delta = 1$ if $S^h(S^g, \Gamma) > 0$.

Suppose for some (X, B^p, Γ, y) the government chooses $\delta = 0$ and S^g, A' such that $S^h(S^g, \Gamma) > 0$. The choice of δ, S^g, A' determines $B', S, B^{p'}$ through the constraints of (14). 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 $S^h(S^g, \Gamma) > 0$.

A.2 Marginal felicity

The second claim is that the marginal felicity for households with respect to S^g strictly increases when spending falls below the threshold for a fiscal emergency while the marginal felicity for the government strictly decreases if $y^h > 0$.

If $y^h = 0$ then households have no income to transfer. The marginal felicity w.r.t. S^g for both households and the government is $u'(S^g)$ which is always strictly decreasing in S^g . The more interesting case is when $y^h > 0$. Holding Γ fixed, let θ denote the threshold for a fiscal emergency, i.e. the lowest value of S^g such that $S^h(S^g, \Gamma) = 0$.

For $S^g \leq \theta$, the transfer function $S^h(S^g, \Gamma)$ solves $u'(S^g + S^h) = v'(Y^h - S^h)$. Thus, we have that for $S^g \leq \theta$

$$S_1^h(S^g, \Gamma) = -\frac{u''(S(S^g, \Gamma))}{u''(S(S^g, \Gamma)) + v''(C(S^g, \Gamma))} \in (-1, 0).$$

If $S^g \leq \theta$, the marginal felicity for households w.r.t. S^g is $u'(S(S^g, \Gamma)) [1 + S_1^h(S^g, \Gamma)] -$

Table A1: Revenue process parameters by city type

	g^{exp}	g^{rec}	ρ	σ_{large}	σ_{small}
Large cities	1.048 (0.002)	0.946 (0.004)	0.406 (0.017)	0.112 (0.004)	
Small cities	1.048 (0.003)	0.942 (0.007)	0.362 (0.017)		0.177 (0.007)

Notes: This table shows the parameters of the city revenue process estimated separately for large and small cities. 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 (19) on a panel of revenue for 2003-2015. The first row shows the results when the regression is run solely using large cities. The second row shows the result when the regression is run solely on small cities. The final two columns give the standard deviation of the idiosyncratic shock, which is calculated from the residuals of the regression.

$v'(C(S^g, \Gamma)) S_1^h(S^g, \Gamma) = u'(S(S^g, \Gamma))$ since the transfers ensure that the marginal felicity of services and consumption are equal. If $S^g > \theta$, then the marginal felicity for households w.r.t. S^g is still $u'(S(S^g, \Gamma))$. Since $S_1^h(S^g, \Gamma) > -1$ for $S^g \leq \theta$, we know that for all S^g , the total value of services $S(S^g, \Gamma)$ is strictly increasing in S^g and $u'(S(S^g, \Gamma))$ is strictly decreasing in S^g . Thus, marginal felicity for households w.r.t. S^g strictly increases when S^g decreases.

For the government, the marginal felicity w.r.t. S^g is $u'(S(S^g, \Gamma))$ when $S^g > \theta$ and $u'(S(S^g, \Gamma)) [1 + S_1^h(S^g, \Gamma)]$ when $S^g \leq \theta$. Since $S_1^h(\theta, \Gamma) < 0$, the marginal felicity for the government discontinuously decreases when spending falls below the fiscal emergency threshold from $u'(S(\theta, \Gamma))$ to $u'(S(\theta, \Gamma)) [1 + S_1^h(\theta, \Gamma)]$.

B Quantification

B.1 Large and small city revenue

Table A1 shows the results when the revenue process parameters g^{exp}, g^{rec}, ρ are estimated separately for large and small cities. For the growth in expansions g^{exp} and recessions g^{rec} , the values are nearly the same for large and small cities and the differences are not statistically significant. For the persistence of the idiosyncratic shocks ρ , the difference between large and small cities is statistically significant but not economically meaningful. While a difference in persistence of 0.04 would be very important if these persistences were near 1, this difference is relatively unsubstantial when both persistences are below 0.5. To summarize, the estimation shows that large and small cities do not differ in their aggregate trends and that idiosyncratic deviations from these

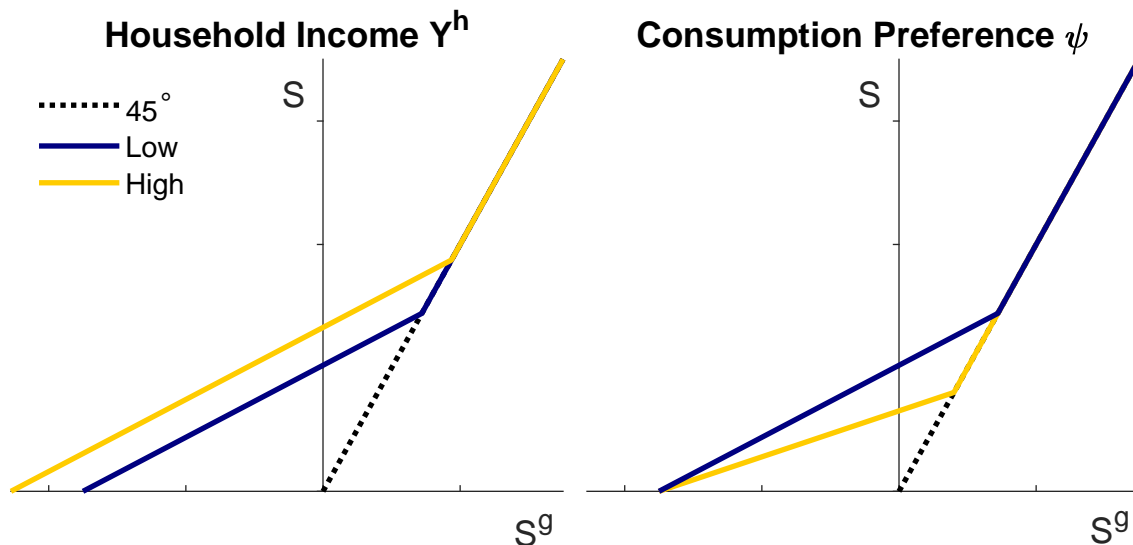


Figure A1: Effect of household parameters on spending function $S(S^g, \Gamma)$

This figure shows total spending as a function of government spending for a fixed Γ . This is calculated for the case where preferences for services and consumption are CRRA, which results in households choosing a linear transfer function. The left panel shows the services as a function of government spending for low (dark blue) and high (light yellow) values of household income. The right panel shows the services as a function of government spending for low (dark blue) and high (light yellow) values of households' consumption preference. In both panels the dotted line shows the 45 degree line where $S = S^g$, i.e. there are no transfers from households.

trends are short-lived.

B.2 Identification

This section discusses how the values for the parameters y^h and $\bar{\psi}$, which control households' income and preference for consumption, are identified by the savings and risk exposure decisions of large and small cities. Given the CRRA preferences specified in Section 4.1, the transfer function for households is

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

From equation (16), this means that services are

$$S(S^g, \Gamma) = \max \left\{ \frac{1}{1 + \psi^{1/\gamma}} (Y^h + S^g), S^g \right\}. \quad (\text{B.2})$$

In other words, total spending is a piecewise linear function of government spending S^g . When $S^g < \psi^{-1/\gamma} Y^h$, the government declares a fiscal emergency and the value of total services S will depend on household income Y^h and households' preference for private consumption ψ .

Importantly, Y^h controls the level of services in a fiscal emergency while ψ controls the slope. Figure A1 shows S as a function of S^g for a fixed Γ in dark blue and how this changes when Y^h or ψ are increased in light yellow. Any differences between S and the 45 degree line are due to transfers from households. How much the government chooses to save will depend mainly on the level of services in fiscal emergencies, i.e. the amount of resources households have available to give to the government. In comparison, the government's choice of risk exposure will depend on the slope of services in fiscal emergencies, as this controls the relative value of services during moderate and severe reductions in S^g .

The parameters that control Y^h, ψ are y^h and $\bar{\psi}$ through equations (5) and (18). As discussed in Section 4.3, the choices of large cities are more affected by the transfer function than the choices of small cities because large cities are more likely to declare a fiscal emergency. Thus, the value of y^h is identified by the amount that large cities save relative to small cities, while $\bar{\psi}$ is identified by the amount of risky assets large cities choose relative to small cities.