
ARE STATE PUBLIC PENSIONS SUSTAINABLE?
Why the Federal Government Should Worry About State Pension Liabilities

Joshua D. Rauh

This paper analyzes the flow of state pension benefit payments relative to asset levels and contributions. Assuming future state contributions fund the full present value of new benefits, many state systems will run out of money in 10-20 years if some attempt is not made to improve the funding of liabilities that have already been accrued. The expected shortfalls raise the possibility that the federal government will be faced with a decision as to whether to bail out states driven to insolvency by their pension programs.

Keywords: public pensions, state budgets, fiscal policy, pension reform

JEL Codes: H55, H72, H74

Joshua D. Rauh: Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL, USA (joshua-rauh@kellogg.northwestern.edu)

I. INTRODUCTION

Just like the federal government, state governments carry substantial amounts of off-balance-sheet debt. At the federal level, the off-balance-sheet liabilities are mostly for programs that cover broad segments of the U.S. population, such as Social Security and Medicare. At the state level, most of the off-balance-sheet debt comes in the form of pension liabilities owed to a narrower group of individuals: current and former public employees. In recent work, Novy-Marx and Rauh (2009a, 2010) have shown that the difference between state public pension liabilities and the assets set aside to fund them is substantially greater than the municipal debt recognized on state balance sheets. Discounting the benefit cash flows at Treasury rates, for example, the gap between assets and already-promised liabilities in state pension funds alone was over \$3 trillion at the end of 2008. This compares to \$1.00 trillion in other forms of recognized state debt under U.S. Census Bureau measures.¹

This analysis raises the question of how long state governments can continue to run their employee pension programs before the pension funds run out of money. The answer of course depends on assumptions about the flow of contributions, benefit payments, and investment returns in the future. In this paper, I assume that state contributions to pension funds will be sufficient to fully fund newly accrued or recognized benefits at state-chosen discount rates (usually 8%), but no more. I show that this assumption is broadly in keeping with states' recent contribution behavior. Using the model of public pension benefit cash flows from Novy-Marx and Rauh (2010) to estimate benefit payments from existing liabilities, the analysis then

¹ See U.S. Census Bureau, 2009, "State Government Finances: 2008,"

<http://www2.census.gov/govs/state/08statess.xls>

examines when state pension funds would in aggregate run out money under different asset return assumptions.

Based on September 2009 asset values, if state pension fund asset returns have an average return of 8% going forward (the states' typical assumption), states in aggregate will run out of funds in 2028. If average returns are 10% through 2045, the funds in aggregate will be roughly sufficient to cover liabilities to existing workers under the states' actuarial assumptions. If average returns are only 6%, state funds in aggregate will run out in 2024. This analysis assumes that state inflation forecasts, which average 3%, are met. If inflation is greater holding the investment outcomes fixed, then even under the higher asset returns the funds will run out sooner, as many state systems provide inflation-linked cost of living adjustments (COLAs) to beneficiaries.

There is substantial cross-sectional variation in the health of the pension plans. Assuming 8% asset returns, Illinois would run out in 2018, followed by Connecticut, New Jersey, and Indiana in 2019. Five states never run out, including New York and Florida, and 17 other states have a horizon of 2030 or beyond. If all states experience 8% average returns, 20 of the states will have run out of pension money by 2025. If the average returns are 10% then only 11 will have run out by 2025. If returns are 6% then 31 will have run out by 2025.

Various other risk factors affect the actual run out dates. The run outs could happen sooner if workers start retiring early in anticipation of problems, if taxpayers start moving out of troubled states, or if contributions are deferred or not made. The run outs could happen later if states make fundamental reforms or can borrow enough to fill the hole. The run outs also would happen later if states use future contributions not to fund new benefits but rather to pay for the

benefits of existing workers, although in that scenario the run outs are more likely to happen as states are digging themselves into a deeper and deeper hole.

If states wanted to remedy this situation over the next 10 years with supplemental contributions, total contributions would have to rise by \$75 billion annually, again assuming 8% investment returns. For comparison, total 2008 state tax revenues were \$781 billion, and annual contributions in 2008 were approximately \$100 billion. Thus, annual contributions would have to rise by 75% during the coming decade.

I also review evidence from Novy-Marx and Rauh (2009b) that the condition of public pension plans has already affected state borrowing costs and is likely to continue to do so. During the 3 months ending December 2008, losses in state pension funds amounted to between 1% and 6% of annual gross state product, and between 9% and 48% of annual state revenue, depending on the state. Tax-adjusted municipal bond spreads rose by 10-20 basis points for each 1% of annual gross state product (or 10% of annual state revenue) lost in pension funds by states in the lower half of the credit quality spectrum. This suggests that if states are planning to roll over unfunded pension liabilities into bonds when pension funds eventually run dry, they may find themselves doing so at substantially higher borrowing costs than they face today.

This paper proceeds as follows. In Section II, I describe how state pension cash flows associated with existing promises are derived. Section III explains how these cash flows inform the question of when states will run out of money. Section IV examines the consequences of this state of affairs, and Section V concludes.

II. Modeling State Pension Cash Flows

From the perspective of an economist studying state pensions, it would be useful if states presented complete forecasts of the stream of cash flows they owe to beneficiaries under

different assumptions. Streams of cash flows can be analyzed in several ways. First, one can examine the determinants of these streams of cash flows to ascertain how certain they are and with what risk factors they covary. Second, they can be compared to existing levels of pension fund assets to determine how long the pension funds are likely to last before the state needs to draw on additional funds. Third, they can be discounted at rates that reflect their risk to arrive at present-value measures of the public liability.

Unfortunately, public pension disclosures do not present forecasts of long-horizon benefit payments. Instead, they present an Accrued Actuarial Liability (AAL), a present value measure of the cash flows under a discount rate chosen by the states to conform with Government Accounting Standards Rule 25 (1994). This rule, which I will refer to as GASB 25, stipulates that states should discount pension obligations at an expected rate of return on pension plan assets. In Novy-Marx and Rauh (2009a), we explain in detail why this approach is misguided. The fundamental problem is that the present value of the liability under this rule depends on the assets the states choose to fund the liability. The riskier the assets the states choose, the higher expected rate of return the assets would have and the smaller the liability would appear. This ignores the fact that riskier assets also have a wider distribution of outcomes and therefore ignores risk completely.

Define $L_{i, \text{stated}}$ as the liabilities that a given state i reports, and define $r_{i, \text{stated}}$ as the flat discount rate that the state reports it uses. States are discounting cash flows using a formula

$$L_{i, \text{stated}} = \sum_{t=1}^T \frac{C_{i,t}}{(1 + r_{i, \text{stated}})^t}$$

but they do not report the cash flows ($C_{i,t}$), which appear in the numerator.

If we want to know how long state pension funds are likely to last, or examine the sensitivity of the liability measurement to state assumptions, it is therefore necessary to back out

the stream of payments which states used as the numerator in their discounting exercise. Once that stream is estimated, it can be analyzed relative to pension fund assets.

Fortunately, state pension systems do provide substantial actuarial and demographic information about their workforces in their annual reports, including benefit formulas, the numbers and average wages of state employees by age and service, salary growth assumptions by age, mortality assumptions, cost of living adjustments (COLAs), and separation (job leaving) probabilities by age. In Novy-Marx and Rauh (2010) we draw on these payments to model a projected stream of payments, our best estimate of the pension-related cash flows that each state used in its discounting procedure to arrive at its stated liability. We collected the data from the Comprehensive Annual Financial Reports (CAFRs) of 115 major pension plans sponsored by the 50 U.S. states.²

A number of subtle issues arise in the exercise of estimating the stream of payments a state has promised to its current and former employees. In a typical defined benefit pension plan, a worker accrues the right to an annual benefit upon retirement that equals a flat percentage of that worker's final (or late-career) salary times the worker's years of service with the employer. Thus, for a given worker, both the years of service and the salary will grow with each year of work, so that the nominal retirement benefit that a worker expects to receive increases more than proportionately with the worker's age.

This raises the question of whether it is sufficient for the state to recognize as a liability only the pension benefit that a given worker has earned based on his work until now. That practice results in a narrow measure. It would be substantially less than the present value of the

² In an earlier version of this study, there were 116 plans. The Illinois Municipal Retirement Fund has been excluded on the grounds that it is funded exclusively at the local level, even though benefits are set by the state.

benefits the state *expects* to pay out, given that most workers are going to continue working and receive salary increases. Broader measures of the liability consider projections of how many more years the state expects the employee to work and at what rate wages will increase. These broader measures, however, would not credit the state in full for the possibility that it could freeze future pension accruals for existing workers.

As an example, consider a worker in a plan with a 2 percent “benefit factor.” The worker has 10 years of job tenure and has an average wage in the last several years of work equal to \$50,000. If this worker quits the job today, the worker would be entitled to a pension of \$10,000 per year ($2\% \text{ per year} \times 10 \text{ years} \times \$50,000$) upon reaching retirement age, plus any cost-of-living adjustments the plan offers. The state could view the liability as an obligation to pay this \$10,000 per year plus cost of living adjustments, but no more. This is the so-called Accumulated Benefit Obligation (ABO) approach. Alternatively, it could recognize the fact that the employee has not yet quit or retired, that she will do so each year with a certain statistical likelihood, and that she will almost certainly be earning more than \$50,000 the year before she retires. The Novy-Marx and Rauh (2010) model allows us to translate liabilities among these different accrual approaches.

Most states report accrued actuarial liabilities under the so-called Entry Age Normal (EAN) method, in which new service liabilities accrue as a fixed percentage of a given worker’s salary throughout his or her career. The EAN is therefore broader than an ABO, in that it recognizes some future pay increases and future service. The EAN measure is not as broad as a full Present Value of Benefits (PVB), which would incorporate all actuarial projections of how long public workers are likely to work. In this paper, I focus on the ABO, as it reflects that benefit cash flows that states have already promised under contractual agreements that are the

least subject to any default, especially given state constitutional protections. Note also that benefits that retirees are already getting and benefits accrued for current service and salary have not been touched in past municipal crises (see Brown and Wilcox (2009)).

Figure 1 shows the stream of payments that the model generates under the three different actuarial methods described above. The model requires the stream of payments to be consistent with the state disclosures: the present value of liabilities, the discount rate, the benefit formulas, and other actuarial variables. The model is calibrated for each state so that the discounted cash flows yield the state's liability $L_{i, stated}$ under the discount rate that the state reports that it is using. Since most state plans use the EAN, discounting these cash flows at state-chosen discount rates reproduces the state's stated liability. The ABO is the lowest of these lines, as it represents only benefit flows that employees have contractually earned based on their service and salary as of 2008. These payments include the expected value of all COLAs, as states include the COLA part of the benefit in calculating the present value of the liability. Figure 1 does not in any way reflect expected payments that will be owed to workers who are not yet hired.

The full details of the model are given in Novy-Marx and Rauh (2010). The calculation essentially begins with a five-item array for each of the 115 plans that includes: 1.) the plan's stated liability; 2.) its state-chosen discount rate; 3.) the actuarial method employed by the state to calculate the liabilities; 4.) the share of active workers in the plan; and 5.) the share of retired workers in the plan, where the remaining group are vested workers no longer in state employment. The conversion of this array into the modeled stream of cash flows requires some assumptions about salary growth by age, separation probabilities by age, the distribution of plan participants by age and years of service (an "age-service matrix"), the average wages of employees in each cell of the age-service matrix, and the average ages and wages of retirees. To

derive these inputs, we examined the financial reports of the 10 states with the largest total liabilities and took the assumptions from the reports where they were usable: New York, Illinois, Pennsylvania, Ohio, and Texas.

The model also requires several additional technical assumptions, specifically: i.) a benefit formula, which we assume to be a constant fixed fraction of salary times years of service, with employees vesting after five years service; ii.) a cost of living adjustment (COLA); iii.) an inflation assumption, which we require to estimate the benefits of annuitants; iv.) a vector of mortality assumptions by age; v.) an age at which employees can retire with full benefits; and vi.) an age at which they can retire with reduced benefits.

For (i)-(iii), we collected data from the individual plans and from the Center for Retirement Research on a plan-by-plan basis.³ For mortality, we use the RP2000 combined mortality table for all plans. For retirement ages, we tested a variety of assumptions and found that assuming public employees can retire at 55 with full benefits and as early as 50 with reduced benefits generated results that best fit the moments of the data. In particular, this assumption allowed the model to closely replicate actual 2008 benefit payments, as well as total liabilities when we ran the model “from scratch” without calibrating to each state’s reported $L_{i, stated}$. Retirement at 55 is consistent with the offerings of many plans, including most programs in the California Public Employee Retirement System (CalPERS), the largest of all the US programs. A number of other systems have later retirement ages, but some such as those designated for public safety officials have even earlier ages. As an early-retirement benefit reduction factor, we use 6% for each year the member is under the early retirement age, which we found to be a common

³ See Center for Retirement Research State and Local Pension Data for 2006, at http://crr.bc.edu/frequently_requested_data/state_and_local_pension_data_4.html

practice, although it is actuarially generous in that the expected present value of future benefits under early retirement is greater than that under normal retirement.

Annual benefit payments under the ABO method are currently around \$150 billion, but they will rise to \$250 billion by 2025. Annual payments under the EAN method would peak slightly later at \$300 billion. The ABO line represents the approximate benefit cash flows that one would expect if all of the plans were frozen at today's level of each worker's wages and years of service. Assuming state pension plans are not frozen, benefit payments will grow substantially more than the ABO, but they will also grow substantially more than the EAN, as the EAN method itself only recognizes part of the actual expected future benefits. In fact, if plan benefits continue to accrue under current formulas and rules, the benefit flow would actually reach \$400 billion in the late 2020s and peak and around \$425 billion in the mid 2030s, as shown by the line labeled PVB.

The analysis that follows in Section III examines how long the assets set aside in pension funds today will last if the ABO cash flows in Figure 1 are the only ones for which the current pension fund assets will be used. Focusing on ABO liabilities assumes that contributions to pension funds going forward will be dedicated and sufficient to meet the present value of all new ABO liabilities. In other words, the analysis is assuming that future contributions to pension funds (net of benefits paid) will exactly meet all changes in the ABO liability that are not due to the fact that each year a given benefit cash flow gets one year closer.

When there is growth in the stated liability not due simply the passage of time, it comes primarily from two effects. First, employees work longer and receive pay increases, none of which is projected into the current liability under the ABO method. These projected future benefit increases are included fully in the PVB, and to some extent in the EAN method. Second,

there are often instances where the true experience of the plan's participants and retirees is different from the actuarial assumptions. For example, workers may work longer than actuaries expected, or live longer, or retire earlier. This kind of difference between experience and actuarial assumption is called an "actuarial loss". The baseline assumption in this paper is that future contributions will meet liability increases due to service and actuarial loss, but they will not help to remedy any shortfall with respect to existing liabilities. Note that if actuaries are incorrect about these assumptions, it will also mean that the EAN and PVB estimates as of today are too low.

Figure 2 examines whether this assumption is reasonable by graphing total contributions against an estimate of new liabilities for 2006-2008. The contributions are taken directly from the annual reports of the 115 major state-sponsored pension systems across the 50 states. The new liabilities are estimated as the change in the ABO liability under state-chosen assumptions, plus benefits paid during the year, minus the increase in the liability due to the passage of time under the state-chosen discount rate (the so-called "interest cost"). This follows directly from considering the increase in the liability over time as deriving from three factors: 1.) the accrual of new benefit promises, which raises the liability; 2.) the payment of existing benefit promises, which lowers the liability; and 3.) the fact that as time passes the cash flows are discounted by one fewer period.

Figure 2 shows that total contributions were \$83 billion in 2006, \$89 billion in 2007, and \$99 billion in 2008. These contributions compare to new benefits in the range of \$88 billion, \$140 billion, and \$99 billion during these three years respectively. Thus, while states contributed exactly the present value of new liabilities in 2008, they contributed less than the present value of new liabilities in 2006 and 2007.

GASB 25 stipulates that states should make annual required contributions (ARCs) which include: 1.) the cost of newly accrued benefits due to service and wage increases; 2.) amortized payments to make up unfunded actuarial liabilities; and 3.) amortized payments to make up any actuarial loss.⁴ If states were making these payments, they would be likely be contributing more than the new liabilities considered above.⁵ However, there is no binding rule or regulation which forces these contributions to be made, and in fact state governments often do not make the full ARC. According to Munnell et al (2008) a full 43% of state and local governments did not pay their ARCs in 2006, and 28% contributed less than 80% of the ARC.

In practice, states may not use future contributions to secure new benefits but rather just use them to pay current benefits. However, Figure 1 shows that new contributions will be needed very soon to meet the newly promised benefits reflected in the PVB cash flows. Also, if public systems find that past actuarial assumptions were too optimistic, those contributions will also be needed to very soon to meet payments not reflected in any of the cash flows in Figure 1. Thus, this practice may buy states a few more years, but if states say they are doing this then they are effectively admitting that they are running an unsustainable scheme by not funding new benefit promises.⁶

In sum, based on Figure 2, the assumption is made that future contributions and future benefits will exactly offset each other. This simplifies the analysis as it allows a consideration of

⁴ Under GASB 25, the maximum amortization period was originally 40 years but was changed to 30 years in 2004.

⁵ It would depend roughly on whether the amortized investment losses were greater than or less than the annual actuarial losses.

⁶ Current employees presumably would like to believe that their contributions are mostly going towards saving for their retirement, rather than paying for current retirees in a system that will have failed by the time they reach retirement.

the existing assets and accrued liabilities in isolation. If ongoing contributions are insufficient to fund new liabilities, then states run the risk of the gap between assets and the present value of liabilities growing even more quickly. If states do not use future contributions to secure new benefits but rather use them to pay current benefits, it may extend the run out period by a few years, but it increases the likelihood of a run out occurring.

III. When Might States Run Out of Money?

As of September 2009, there were roughly \$2.2 trillion of assets in the 115 defined benefit pension plans examined in this study. Given the stream of payments that states are obligated to make, how long are these assets likely to last? In this section, I model the trajectory of aggregate state pension fund asset levels under some calibrated assumptions about public pension fund asset performance, as well as some scenarios in which states make remedial contributions to improve the funding status of already-accrued benefits.

The top graph in Figure 3 shows the baseline analysis. Asset levels in billions of dollars are presented under three assumptions: a 6% annualized return, an 8% annualized return, and a 10% annualized return. States invest in portfolios which according to their calculations have an 8% expected return, so the middle line can be thought of as the average outcome. The evolution of assets is calculated as follows: assets at time $t+1$ are equal to assets at time t times the assumed returns, minus the benefit payments shown in Figure 1. I begin with an aggregate analysis of all 115 plans, and then consider state heterogeneity.

Annualized returns could of course be *substantially* different from 8%. The actual evolution of assets is shown in the figure for 2005-2008 and is clearly volatile. Novy-Marx and Rauh (2009a) calculate that the annual volatility of the portfolios of state pension fund investments is itself on the order of 8%, based on the historical returns of the state pension fund

asset classes. In this paper I consider a substantially narrower band of outcomes. Part of what drives the historical volatility of asset returns is variation in inflation, and the cash flows measured in the section above rely on the state's inflation forecast being accurate. It is appropriate to consider these narrower bands if high inflation is likely to be correlated with higher nominal asset returns, as both asset returns and liability cash flows would be higher at the same time. Furthermore, if there is mean reversion in asset returns, the long-horizon volatility would be much smaller than the short-horizon volatility. Nonetheless, it is worth keeping in mind that substantially more extreme outcomes are possible than the ones presented here.

The top graph of Figure 3 also assumes that contributions in the future are exactly sufficient to fund new liabilities. The exercise therefore addresses the question of how long assets will last if set strictly to the purpose of funding already-recognized pension promises. The graph shows that if returns average only 8%, the existing funds will in aggregate run out in 2028. If returns average 6%, they will run out in 2024, and if returns average 10% they will extend through the 2040s. It is important to interpret these findings in the context of Figure 1, which shows that a large part of already-promised payments are owed in the 2020s and 2030s. Under these assumptions, states therefore will have to raise substantial amounts of new finance in the future just to fund benefit promises that they have already made.

Figure 3 rather unrealistically assumes that pension funds in different states can and will subsidize each other. Table 1 disaggregates the analysis and shows how the projected day of reckoning varies by state. Under 8% realized returns, seven states run out of money before the end of 2020, including Illinois (2018), New Jersey (2019), and Connecticut (2019). 20 states have run out by the end of 2025, and 31 states by the end of 2030.

The damage inflicted by this problem depends upon how large the benefits owed to workers actually are relative to the state's revenues. In Illinois, obligations already recognized today under the ABO will result in \$11 billion in average annual pension payments in 2019-2023, the five years after the funds run dry. Projected payments (i.e. those under the PVB) will be \$14 billion during those years. Tax revenues for the state of Illinois were \$31.9 billion in 2008, according to recent U.S. Census Bureau measures. Moving to a pay-as-you-go system would therefore be a catastrophic shock to the revenue needs of the state of Illinois. As shown in the second column from the right, the benefits in the five years following run-out are 35% of total annual state revenue flows, considering only ABO benefits promised as of 2008 and benchmarking to 2008 revenues. The final column assumes 3% revenue growth and shows that projected annual benefit payments for Illinois will be 32% of projected 2019 revenues.

Some states whose funds might not run out of money until the mid 2020s face a very challenging situation when they eventually do. Ohio collected \$26.4 billion in tax revenues 2008. If their pension funds run dry in 2030, they will face \$13.8 billion of annual benefit payments during 2031-2035, which will have to be made out of general revenues. Those benefit payments represent 52% of 2008 tax revenues, and 55% of projected 2031 revenues under 3% revenue growth.

Across all 45 states that are projected to eventually run out, the mean ratio of already-promised benefits after the run out to 2008 revenues is 28%, and the mean ratio of projected benefits to revenues projected forward at a 3% growth rate is 27%.

States could attempt to remedy this problem earlier rather than later by establishing supplementary contribution programs. The middle graph of Figure 3 assumes that states in aggregate contribute \$50 billion of additional catch-up contributions to their pension funds each

year during 2010-2020. Compared to \$99 billion in actual 2008 contributions, this represents more than a 50% increase. The figure shows that if asset returns average are 8% annually, the funds will be solvent through 2037, 9 years longer than without the \$50 billion of additional contributions. If returns are 6%, the funds will only last through 2028.

If returns are 10% or higher under the scenario of \$50 billion increases in contributions for 10 years, states would actually be in a position where the pension funds are overfunded and money might need to be returned to taxpayers. An overfunding situation poses the risk that the excess funds might be spent inefficiently, particularly if public employees use overfunding as a bargaining tool to prevent wage increases. Nevertheless, the states would have the option of reducing the supplemental contributions during the middle of the upcoming decade if returns had been particularly strong. It is also worth noting that 10% returns over a period of decades with 3% inflation is quite unlikely.

The bottom graph considers a more radical approach which would be for states to contribute an additional \$75 billion per year to pension funds. This would entail a 75% increase in contributions and represents about 10% of total state tax revenues for 2008. Under the baseline asset return assumption of 8%, these contributions would extend the solvency of the state funds through 2045.

In sum, under the baseline assumption of 8% annual asset returns, states need to contribute at least \$75 billion each year over the next 10 years in order to avoid a significant chance of having to draw on resources outside their pension funds just to meet benefits that have already been promised. This entire analysis of course also assumes that contributions going forward are sufficient to pay for all newly promised benefits, an assumption which may be hard to meet given the substantial difference between the ABO and the PVB,

IV. What Are the Consequences?

Ultimately the question of whether it is appropriate for states to have trillions of dollars of off-balance-sheet debt boils down to two separate questions. First, what is the appropriate amount of total state indebtedness? Second, what is the optimal distribution of this indebtedness between on-balance-sheet and off-balance-sheet liabilities?

There is a large class of theories that argues that the total amount of public debt is irrelevant. In terms of the intergenerational consequences of state debt, a starting point is the famous doctrine of Ricardo (1820), which postulates the irrelevance for public welfare of financing current spending with debt versus taxes. One statement that can be made with certainty is if Ricardian Equivalence does not hold, then states should not be as indebted as they might optimally be if it did hold. If citizens do not save today in anticipation of future taxes, then future generations have no defense against the profligacy of today's state governments.

There is also no lack of theories on the consequences of public indebtedness, especially as pertains to its crowd-out effect on the capital stock and the upward pressure it could exert on macroeconomic interest rates (Modigliani (1961), Feldstein (1974)). At the same time, the empirical evidence for these channels is not particularly conclusive, at least in the relatively short time series at the disposal of researchers. These effects are of course muted in open economies. They are perhaps best thought of as the effects on global interest rates and capital if the governments of the world in aggregate borrow more from the non-government sector.

One consequence of public debt that has been explored a little more is the "sovereign default" channel. If international markets lose faith in sovereign debtors, it becomes very expensive for the sovereign to roll over debt. This raises the question of how costly it is likely to be for states to roll over their pension debts into other forms of debt when eventually the pension

funds run out of money. Novy-Marx and Rauh (2009b) examine the effects of losses in U.S. state pension funds during 2008 on state borrowing costs. We found that tax-adjusted municipal bond spreads rose by 10-20 basis points for each 1% of annual gross state product (or 10% of annual state revenue) lost in pension funds by states in the lower half of the credit quality spectrum (Novy-Marx and Rauh (2009b)).

Thus, the state of public pension plans has already affected state borrowing costs and is likely to continue to do so. By the time the public pension debt needs to be rolled over into other forms of debt, states could find their fiscal condition makes such refinancing very expensive. Under current law, such rollovers would also be tax-disadvantaged. Attempts by states in the early 1980s to issue tax-preferred bonds in order to fund state pensions led to IRS rulings deeming these instruments “arbitrage bonds” and disallowing tax deductibility.⁷ Despite the fact that Pension Obligation Bonds (POBs) today enjoy no tax subsidies, they are becoming increasingly popular as states struggle to make contributions, as evidenced by Illinois’s \$3.5 billion issue in early 2010.

In the context of the question of whether it is optimal to have such a large portion of debt be off-balance-sheet, it must be recognized that in many states, constitutional provisions limit the extent of state general obligation debt. For many state governments, therefore, pension underfunding may be an important source of public sector borrowing (Novy-Marx and Rauh (2009a)). On the other hand, if taxpayers are unaware of the extent of the true indebtedness because it is off the balance sheet, they may be less likely to adjust their own savings decisions

⁷ At the time, state borrowing rates in the tax-exempt market were lower than Treasury yields, creating an actual arbitrage. In today’s markets that is no longer the case.

accordingly. Ricardian Equivalence is even less likely to operate, and intergenerational transfers are even larger.

V. Conclusions

Whether public pension promises can be met without drawing substantially on other taxpayer resources depends on several factors. A combination of high asset returns and low inflation would certainly help states meet obligations from their current liabilities. On the other hand, states face the risk that higher inflation and low asset returns could make their systems even more vulnerable. Furthermore, the less desirable outcomes would tend to occur in states of the world where additional revenue generation is costly. State governments face a choice between taking more risk today and funding the liabilities to a greater extent.

State governments also face a difficult choice as pertains to future benefit accruals. The main analysis of this paper assumes that states will use future contributions to fully fund new benefit accruals. It then measures the dates at which existing assets are insufficient to fund already-accumulated benefits. In recent years, total contributions across all 50 states of approximately \$100 billion have been sufficient for funding newly accrued benefits. Supplemental contributions across all 50 states of an additional \$75 billion would extend the funds through most of period during which retirees who have already accrued benefits will be receiving payments.

Alternatively, states could attempt to stop future benefit accruals by freezing benefits at their current levels, re-directing the savings to ameliorating the funding deficit of already-accrued benefits. To do this, there would presumably need to be a compensating differential to induce public employees to accept such a deal and to provide for their retirement security.

Employees would need a defined contribution plan, and those that do not have Social Security (some 25% of state and local workers) would need to be brought into that program.

It is worth noting that in some other countries facing pension shortfalls, the system is set to adjust payments to all program participants including existing retirees. For example, the Netherlands collective labor-related plans have conditional indexation, in which retirees only get cost of living adjustments if the value of assets is adequate (Ponds and van Riel (2009)). The US states face a difficult situation given the hardness of the claims and the specific constitutional protections of accrued pension promises by state governments.

There seems to be a high likelihood that future generations will have to bear the substantial burden of making up pension benefits for previous generations of state employees. While citizens of states that are particularly hard-hit by the pension crisis may be able to escape to other states, an acceleration of this demographic phenomenon would leave a dwindling taxpayer base behind in the states facing the largest liabilities. This would increase the likelihood of a federal taxpayer bailout in which taxpayers in all states would bear the burden of the states in default. The problem of state and local pension liabilities is therefore a problem for all US taxpayers, not just those in the states with the largest deficits.

ACKNOWLEDGEMENTS

I thank the organizers and participants of “Train Wreck: A Conference on America’s Looming Fiscal Crisis,” a conference on January 15, 2010 sponsored by Urban-Brookings Tax Policy Center and the USC-Caltech Center for the Study of Law and Politics. I am grateful to Robert Novy-Marx for countless conversations, to John G. Matsusaka for his discussion, and to Roman Feldman for research assistance.

REFERENCES

Brown, Jeffrey, and David Wilcox, 2009.

“Discounting State and Local Pension Liabilities.” *American Economic Review* 99(2), 538-542.

Feldstein, Martin S., 1974.

“Social Security, Induced Retirement, and Aggregate Capital Accumulation.” *Journal of Political Economy* 82, 905-926.

Krishnamurthy, Arvind, and Annette Vissing-Jorgensen, 2008.

“The Aggregate Demand for Treasury Debt.” Kellogg School of Management Working Paper.

Modigliani, Franco, 1961.

“Long-run Implications of Alternative Fiscal Policies and the Burden of the National Debt.” *Economic Journal* 71, 730-755.

Munnell, Alicia, Kelly Haverstick, Steven Sass, and Jean-Pierre Aubry, 2008.

“The Miracle of Pension Funding by State and Local Pension Plans.” Center for Retirement Research, Issue in Brief #5.

Novy-Marx, Robert, and Joshua Rauh, 2010.

“Public Pension Promises: How Big Are They and What Are they Worth?” University of Chicago Working Paper

Novy-Marx, Robert, and Joshua Rauh, 2009a.

“The Liabilities and Risks of State-Sponsored Pension Plans.” *Journal of Economics Perspectives* 23(4), 191-210.

Novy-Marx, Robert, and Joshua Rauh, 2009b.

“Fiscal Imbalances and Borrowing Costs: Evidence from State Investment Losses.” Kellogg School of Management Working Paper.

Ponds, Eduard, and Bart van Riel, 2009.

“Sharing Risk: The Netherlands New Approach to Pensions.” *Journal of Pension Economics and Finance* 8(1), 91-105.

Ricardo, David, 1820.

Essay on the Funding System, Various editions.

Figure 1

Estimated Payments Owed for Benefits Earned through 2008

This figure shows projected aggregate cash flows by state governments due to public pension promises, as would be recognized under several different actuarial methods. As in Novy-Marx and Rauh (2010), cash flow projections for each state plan are made so that the state plan's reported liability equals the discounted value of these cash flows under the state's chosen accrual method (usually Entry Age Normal) and reported discount rate. Cash flows are then adjusted so that only current salary and years of service are recognized. See Novy-Marx and Rauh (2010) for details.

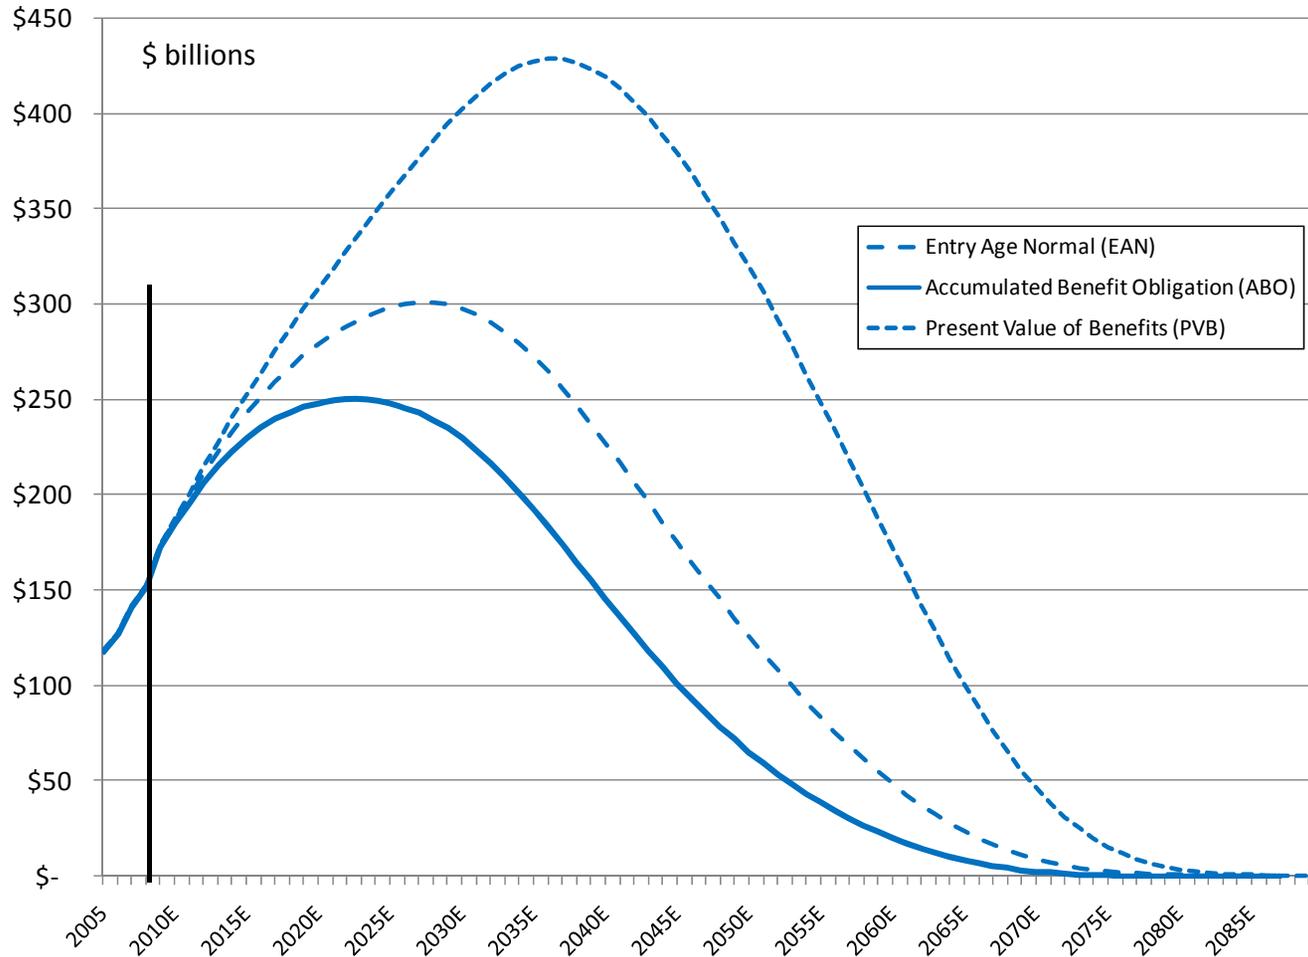


Figure 2

Contributions and New Benefit Accruals in State Pension Plans

The bars show aggregate contributions to the 115 major U.S. state-sponsored employee pension plans, for the years 2006-2008. Contributions come from employers (the white bars), employees (the light gray bars), and other sources such as state appropriations. The line in the graph shows an estimate of new liabilities, equal to the estimated change in the Accumulated Benefit Obligation (ABO) liability plus benefits paid during the year, minus the increase in the liability due to the passage of time (i.e. minus the state-chosen discount rate times the previous year's liability). The new liabilities therefore consist of service costs plus increases that are due to the difference between actual experience and the state's actuarial model.

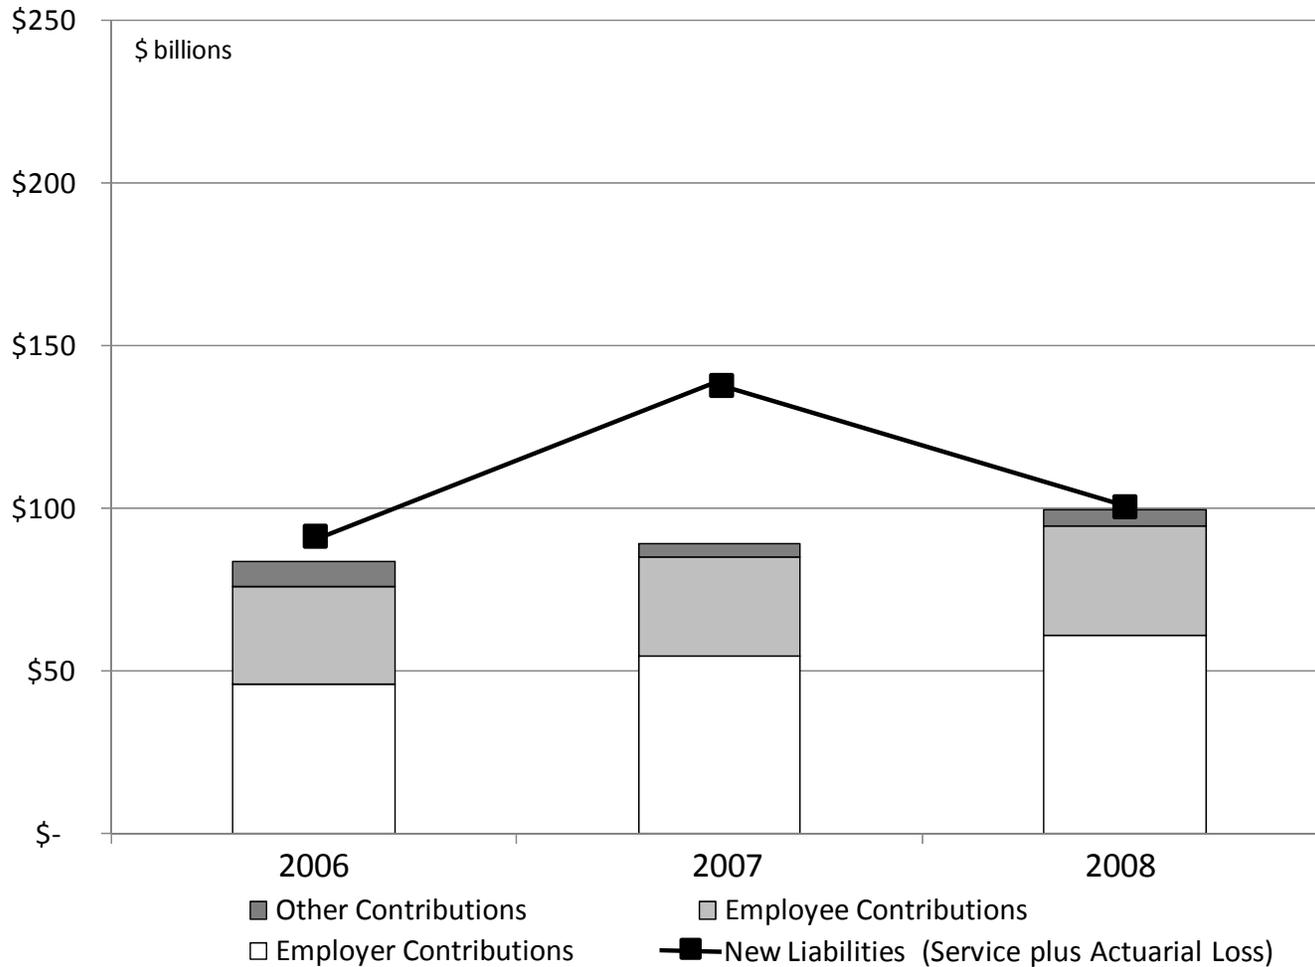


Figure 3

Aggregate State Pension Fund Asset Levels Under Different Scenarios

The three lines in each graph show aggregate state pension fund assets under three return scenarios: a 6% return (the dashed line), an 8% return (the solid line) and a 10% return (the dotted line). The top graph shows assets under the assumptions that contributions only fund new liabilities. The middle graph shows assets under the assumption that states contribute an additional \$50 billion per year above and beyond new liabilities during 2010-2020, and the bottom graph shows assets if states contribute an additional \$75 billion per year above and beyond new liabilities during 2010-2020.

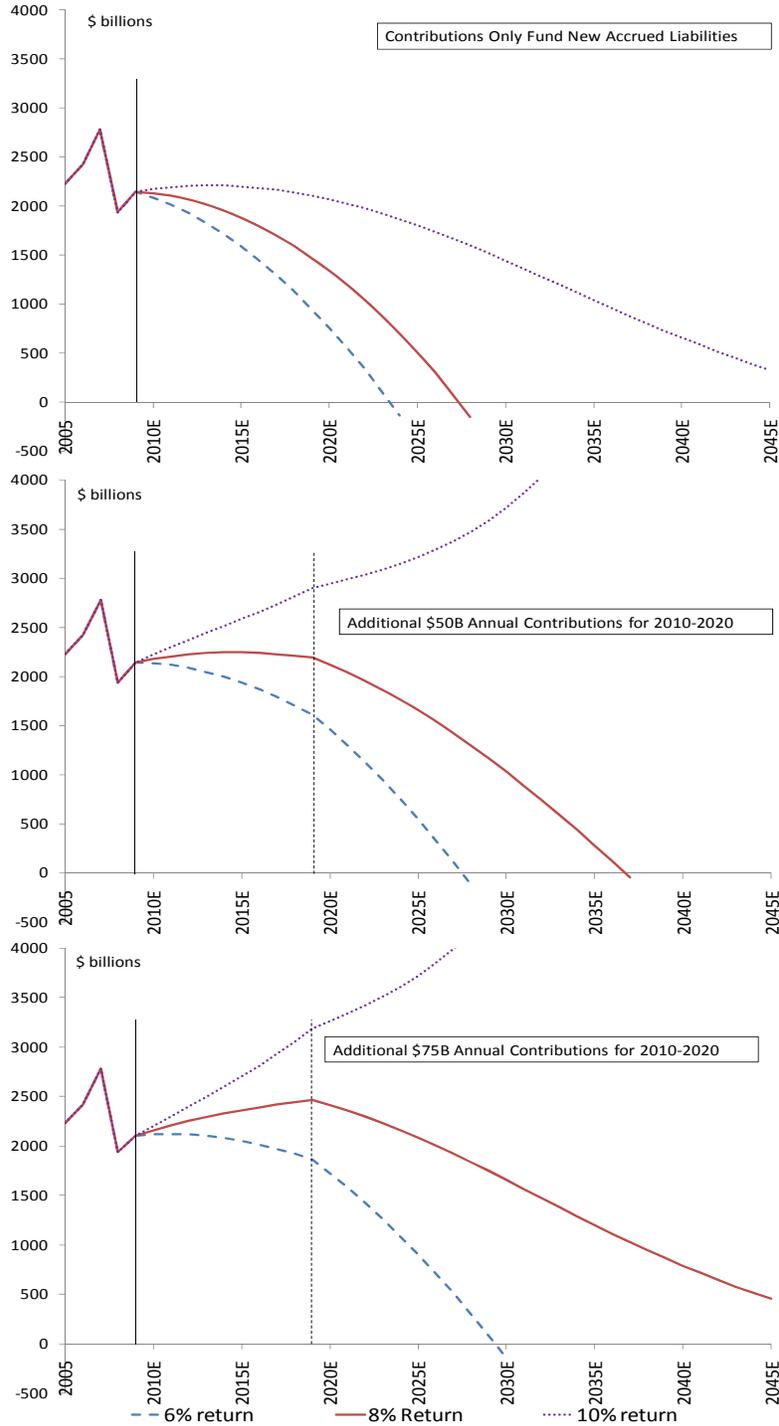


Table 1**When Might State Pension Funds Run Dry?**

The first column is the year that each state will run out of pension fund assets assuming they earn 8% on their assets and that they use future contributions to fund new benefits in full. The starting points for pension assets are the September 2009 values from *Pensions and Investments*. The second column shows the average annual benefit payments that will be owed over the 5 years following the state's estimated run-out, but only including 2008 promises. The third column shows the average annual benefit payments that will be owed over the 5 years following the state's estimated run-out, including projections of future benefits for existing workers. The fourth column shows 2008 tax revenues. The fifth column shows the ratio of annual promised benefits to 2008 tax revenues. The final column shows the ratio of annual projected benefits to forecast tax revenues assuming 3% annual revenue growth. That is, revenues in the year of the run out are estimated by [2008 Tax Revenues] * (1.03)^[Years Left].

year	Year Run Out (8% Returns)	Average Benefit in Five Years Following Runout, \$ billions		2008 Tax Revenues (\$billions)	Ratio of Annual Benefits to Revenues	
		Promised by 2008	Projected as of 2008		Promised, Using 2008 Revenues	Projected, Assuming 3% Revenue Growth
ILLINOIS	2018	11.0	13.6	31.9	35%	32%
CONNECTICUT	2019	3.9	4.9	13.4	29%	27%
INDIANA	2019	3.0	3.6	14.9	20%	17%
NEW JERSEY	2019	10.7	14.4	30.6	35%	34%
HAWAII	2020	1.4	1.7	5.1	26%	24%
LOUISIANA	2020	3.2	4.3	11.0	29%	27%
OKLAHOMA	2020	2.7	3.7	8.5	32%	30%
COLORADO	2022	5.2	7.8	9.6	54%	54%
KANSAS	2022	1.7	2.5	7.2	24%	23%
KENTUCKY	2022	3.7	5.3	10.1	36%	35%
NEW HAMPSHIRE	2022	0.7	1.0	2.3	30%	30%
ALABAMA	2023	3.7	5.5	9.1	41%	39%
MICHIGAN	2023	5.5	7.8	24.8	22%	20%
MINNESOTA	2023	5.0	7.3	18.3	27%	25%
MISSISSIPPI	2023	2.5	3.9	6.8	37%	37%
MARYLAND	2024	4.0	6.0	15.7	26%	24%
PENNSYLVANIA	2024	9.2	13.8	32.1	29%	27%
SOUTH CAROLINA	2024	3.3	4.7	8.5	39%	34%
WEST VIRGINIA	2024	0.9	1.3	4.9	18%	16%
MISSOURI	2025	4.2	6.9	11.0	38%	38%
MAINE	2026	1.1	1.7	3.7	30%	28%
MASSACHUSETTS	2026	4.2	6.6	21.9	19%	18%
NEW MEXICO	2026	2.1	3.9	5.6	38%	40%
MONTANA	2027	0.6	1.1	2.5	26%	25%
RHODE ISLAND	2027	2.4	2.4	2.8	87%	50%
VERMONT	2028	0.3	0.5	2.5	11%	11%
ARIZONA	2029	4.3	7.3	13.7	32%	29%

ARKANSAS	2030	1.5	3.1	7.5	20%	21%
CALIFORNIA	2030	32.4	68.1	117.4	28%	30%
OHIO	2030	13.8	28.0	26.4	52%	55%
WYOMING	2030	0.5	1.0	2.2	23%	25%
SOUTH DAKOTA	2031	0.5	1.0	1.3	40%	38%
NEBRASKA	2032	0.5	1.2	4.2	13%	15%
VIRGINIA	2033	3.4	8.6	18.4	18%	22%
WASHINGTON	2033	3.4	8.0	17.9	19%	21%
DELAWARE	2035	0.4	1.0	2.9	13%	15%
IOWA	2035	1.3	3.0	6.9	18%	20%
TENNESSEE	2035	1.7	4.2	11.5	14%	16%
UTAH	2036	1.3	3.5	5.9	23%	26%
TEXAS	2037	9.9	30.4	44.7	22%	29%
WISCONSIN	2038	4.1	9.7	15.1	27%	27%
OREGON	2039	1.6	6.0	7.3	21%	33%
NORTH DAKOTA	2041	0.2	0.5	2.3	7%	9%
IDAHO	2043	0.4	1.7	3.7	11%	16%
GEORGIA	2047	2.6	9.8	18.2	14%	17%
ALASKA	–	–	–	8.4	–	–
FLORIDA	–	–	–	35.8	–	–
NEVADA	–	–	–	6.1	–	–
NEW YORK	–	–	–	65.4	–	–
NORTH CAROLINA	–	–	–	22.8	–	–
Mean	2028	4.0	–	15.6	28%	27%
Median	2026	2.7	–	9.3	26%	27%
Standard Deviation	7	5.3	–	19.3	14%	10%