Are State Public Pensions Sustainable?*

Joshua D. Rauh
Kellogg School of Management and NBER

December 31, 2009

Abstract

This paper examines the sustainability of public pension systems by analyzing the flow of benefit payments relative to asset levels and contribution rates. Even if states contribute the full present value of new benefit accruals, state pension funds will on average 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. I also review the evidence showing that under a variety of possible market-based discount rates, the gap between assets and accrued liabilities in state public pension plans is substantially larger than the volume of total outstanding state debt. For example, when the already-promised benefit cash flows are discounted at Treasury rates, the gap between assets and liabilities amounts to 25% of outstanding U.S. federal debt. This gap has affected state borrowing costs and is likely to continue to do so.

* Rauh: (847) 491-4462, joshua-rauh@kellogg.northwestern.edu. This paper was prepared for "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 the conference organizers, to Robert Novy-Marx for countless discussions, and to Roman Feldman for research assistance.

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, 2009b) show 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 accrued liabilities in state pension funds alone is $3.25 trillion. This compares to $1.00 trillion in other forms of recognized state debt (see U.S. Census Bureau (2009)).

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 states will make just enough contributions to their pension funds to pay the present value of newly accrued benefits. 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 (2009a) to estimate benefit payments from existing liabilities, this paper's analysis examines when state pension funds will in aggregate run out money under different asset return assumptions.

If state pension fund asset returns have a geometric mean return of 8% going forward (the states' typical assumption), states in aggregate will run out of funds in 2025. If average returns are 10%, the funds will extend to 2030. If average returns are only 6%, the funds will run out in 2021. This analysis assumes that state inflation forecasts of 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 contain inflation-linked cost of living adjustments (COLAs).

Much of the analysis in this paper treats the state employee pension plans in aggregate. However, there is substantial cross-sectional variation in the health of the pension plans. As a result, if all states

experience 8% average returns, 28 of the states' plans will have run out of money by 2025. If the average returns are 10% then only 16 will have run out by 2025. If returns are 6% then 38 will have run out by 2025.

If states wanted to remedy this situation over the next 10 years with supplemental contributions, total contributions would have to rise by more than $100 billion annually even if states achieve 8% annual 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 more than double during the coming decade.

As explained in Novy-Marx and Rauh (2009a), the benefit payments that underlie these calculations represent liabilities for the states. An expanded state balance sheet that included the financial impact of public employee pension systems would include the present discounted value of these liabilities as state debt, just as it would include the value of pension assets as state assets. The choice of discount rate depends on the perspective that such a balance sheet takes and on the risk of the benefits. If the state is assessing the pension benefits as promises, it would need to recognize the cost of buying financial securities that would provide a default-free stream of payments to beneficiaries. In that case, it is clear that a yield curve free of default risk should be used. Alternatively, one could imagine assessing how much the liabilities are actually worth allowing the state the possibility that it might default on them. In that framework, it might be appropriate to consider a (taxable) municipal yield curve. Even then, however, the state's ability to default is limited by state constitutional provisions that give priority to state employee pension benefits (see Brown and Wilcox (2009)).

I review these discounting arguments and calculations and their implications for the state's effective balance sheet (Novy-Marx and Rauh (2009a, 2009b)). Counter to the logic of financial economics, the discounting practice required by the Government Accounting Standards Board (GASB) entails using a discount rate based on the assumed rate of return on plan assets. In practice, actuaries implement this as a flat 8% discount rate for all horizons on the yield curve. This practice understates liabilities and the funding status of public pension plans rather substantially.

I also review evidence from Novy-Marx and Rauh (2009c) that the state 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. I review evidence 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. 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.

1. 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. 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 risks 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 (2009b), 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_{stated} as the liabilities that a given state reports and r_{stated} as the flat discount rate that the state uses, which is also reported. States are discounting cash flows as follows

$$L_{stated} = \sum_{t=1}^{T} \frac{C_t}{(1 + r_{stated})^t}$$

but they do not report C_t, the numerator cash flows in the discounting formula.

If we want to know how long state pension funds are likely to last, as well as the true risk-adjusted present value of the cash flows, 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 and also discounted at rates that reflect the covariance of the payments with pricing factors.

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 (2009a) 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 116 major pension plans sponsored by the 50 U.S. states.

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. To do so would result in a relatively narrow measure. It would be substantially less than the present value of the 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 would 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 for its ability to freeze future pension accruals.

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% per year × 10 years × $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 when she retires. The Novy-Marx and Rauh (2009a) model allows us to translate liabilities among these different accrual approaches. Indeed, the CAFRs typically report the liabilities on a somewhat broader basis, and we harmonize them all to the ABO (see Novy-Marx and Rauh (2009a) for details).

Figure 1 shows the stream of ABO-related payments that our model generates. 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 figure shows the total payments that we project states will owe to beneficiaries by year based only on liabilities that have been accrued through 2008. These expected payments include the expected value of all COLAs, as states are instructed to include the COLA part of the benefit in calculating the present value of the liability. The figure does not consider benefits that will be earned in the future, either by current workers or workers yet to be hired,

nor does it reflect anything about future contributions. Annual benefit payments are currently around $150 billion, but they will double to $300 billion by 2025.

Assuming state pension plans are not immediately frozen at current levels of accrued liabilities, benefit payments will grow substantially more than shown in the figure. The figure, and this paper in general, focuses only on ABO cash flows. The analysis that follows in Section 2 examines how long the assets set aside in pension funds today will last if the cash flows in Figure 1 are the only ones for which the pension fund resource will be used. Focusing on ABO liabilities is equivalent to assuming that contributions to pension funds going forward under current policy will be sufficient to fund the present value future service-related benefit accruals, but no more. In other words, in the absence of policy changes, future contributions to pension funds will exactly meet increases in the ABO that are due to employees working more years and receiving pay raises. The baseline assumption implies that future contributions will not help to remedy any shortfall with respect to existing liabilities.

Figure 2 examines whether this assumption is reasonable by graphing total contributions against the service cost component of new liabilities for 2006-2008. The contributions are taken directly from the annual reports of the 116 major state-sponsored pension systems across the 50 states. The service cost part of new liabilities is estimated as the change in the ABO liability at the state-chosen discount rate, plus benefits paid during the year, minus the increase in the liability due to the passage of time (the so-called “interest cost”). This follows directly from considering the increase in the ABO over time as deriving from three factors: 1.) the accrual of new benefit promises (which raises the ABO); 2.) the payment of existing benefit promises (which lowers the ABO); 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 $100 billion in 2008. These contributions compare to new benefit accruals in the range of $90-$140 billion annually during this time period. Thus, in the recent past, it has been the case that states have contributed somewhat less than the present value of total new benefit accruals.

GASB 25 does stipulate that states should make annual required contributions (ARCs) which include the cost of newly accrued benefits *plus* amortized payments to make up unfunded actuarial liabilities. 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. This suggests that future benefit accruals are unlikely to be fully funded and that the expected funding deficit is even worse than our findings suggest.

In sum, based on Figure 2, the somewhat generous assumption can be made that future contributions and future benefit accruals will exactly offset each other. This simplifies the analysis as it allows a consideration of the existing assets and accrued liabilities in isolation. However, it is worth noting that if states do not fund the present value of newly accrued liabilities with their ongoing contributions, they run the risk of the gap between assets and the present value of liabilities growing even more.

2. When Will States Run Out of Money?

Given the stream of payments that states are obligated to make, how long is the roughly $2 trillion in assets likely to last? The answer of course depends on assumptions about the flow of investment returns in the future. 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. Assets 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 2.

Annualized returns could of course be substantially higher than 10% or lower than 6%. The actual evolution of assets is shown for 2005-2008 and is clearly volatile. Novy-Marx and Rauh (2009b) 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 is 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 also assumes, as in the previous section, 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-accumulated pension promises. The graph shows that if returns average only 8%, the existing funds will suffice through 2024. If returns average 6%, they will run out before the end of 2022, and if returns average 10% they will extend to 2030. It is important to interpret these findings in the context of Figure 2, which shows that a large part of future 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.

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 $100 billion in actual 2008 contributions, this represents a 50% increase. The figure shows that if asset returns average 10%, these supplemental contributions would be sufficient to keep the pension funds solvent through 2045. However, if returns are 8%, the funds will only be solvent through

2028, just 4 years longer than without the $50 billion of additional contributions. If returns are 6%, the funds will still only last through 2024.

The bottom graph considers a more radical approach which would be for states to contribute an additional $100 billion per year to pension funds. This would entail a doubling of contributions and represents 13% 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 2036. If returns are 10% or higher, states would actually be in a position where the pension funds are overfunded and the money would optimally 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.

The above analysis considers the state employee pension plans in aggregate. However, there is substantial cross-sectional variation in the health of the pension plans. As a result, if all states experience 8% average returns, 28 of the states' plans will have run out of money by 2025. If the average returns are 10% then only 16 will have run out by 2025. If returns are 6% then 38 will have run out, by 2025.

In sum, under the baseline assumption of 8% annual asset returns, states need to contribute at least $100 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 to meet accrued benefits.

3. What Are These Liabilities Worth?

This section draws heavily on Novy-Marx and Rauh (2009a). As discussed in that paper, a financial stream of payments should be discounted at a rate that reflects its risk (Modigliani and Miller (1958)), and in particular its covariance with priced risks (Treynor (1961), Sharpe (1964), Lintner (1965)). We analyze this question in the following way. Suppose that a state wanted to defease its entire pension obligation by paying off the beneficiaries with a portfolio of bonds that generates the same stream of payments as the benefits and defaults in exactly the same states of the world. The cost of the defeasance would be the market value of the bonds the state must deliver today.

If the ABO liability were purely risk-free, it would be appropriate to discount the cash flows represented Figure 1 using a truly risk-free term structure of interest rates. In reality, there are at least four different sources of risk that we might want to consider: i.) default risk, ii.) inflation risk, iii.) liquidity risk, and iv.) wage risk. It is important to emphasize that in the context of an asset-pricing approach to this valuation, these risks only matter insofar as they covary with pricing factors such as consumption or the stock market. Default risk is the risk that the state might simply not make the promised payments. The greater the default risk, the higher the appropriate interest rate for determining the value of the ABO liability. Inflation risk is the risk inherent in any non-indexed bond-like stream of cash flows that the spending power of the payments may be reduced by the time the payments are made. Non-indexed pension payments would face inflation risk and their expected cash flows should be discounted at higher interest rates than COLA-indexed pension payments, other things equal. Liquidity risk refers to the fact that if these claims were tradable, they would not be as liquid as, say, Treasury bonds. Holding other factors constant, the relative illiquidity of a financial claim on pension payments would raise the interest rate. Wage risk is the risk that because of the evolution of future wages, payments may be higher or lower than expected.

Brown and Wilcox (2009) argue that benefits related to the ABO liability are quite safe and should be discounted at something approaching a risk-free rate. Recall that benefits associated with the ABO liability consist of benefits being paid to current retirees, benefits owed to people who are not yet retired but who have left public employment, and benefits already accrued by active workers under benefit formulas. Let us consider how the four potential risk factors might be relevant for this liability.

Consider first default risk (risk (i) above). If government pension liabilities default in the same states of the world as other government debt, and have the same recovery rates, then the discount rate should bear some relation to municipal bond yields, which vary with the credit quality of the issuing entity. The main difference would be that pension benefits are taxed. However, Brown and Wilcox (2009) document that in the majority of the 50 U.S. states, public pension obligations are specially protected by state constitutions in ways that make membership in a pension plan an “enforceable contractual

relationship." These provisions seem to give special protection to accumulated pension benefits. It seems that the risk of state default on the benefits related to the ABO liability is very small.

Consider next inflation risk (ii), liquidity risk (iii), and wage risk (iv). If benefits were not inflation-adjusted, it would make sense to discount the streams of benefit payments at rates similar to those on other nominal bonds. The fact that benefits typically have COLAs suggests that the optimal discount rate should remove the inflation yield premium that raises bond yields higher than what they would be in the absence of inflation risk. On the other hand, pension benefits are less liquid than bonds, suggesting that the optimal discount rate should remove the liquidity price premium that lowers bond yields below what they would be in the absence of this liquidity premium. As for wage risk (iv), since the components of the ABO are invariant to the evolution of future wages, the cash flows related to the ABO liability are not subject to any correlation between real wages and the stock market.

In Novy-Marx and Rauh (2009a), we present liabilities under three main discount rates: 1.) the state-chosen rate, which always has a flat term structure; 2.) a taxable muni yield curve, which is the yield curve of the state-specific credit rating grossed up by 25% to reflect the tax preference; and 3.) the Treasury yield curve. Table 1 reviews the results of the discounting. Under state-chosen rates, total ABO liabilities are $2.80 trillion and the gap between liabilities and assets is $0.86 trillion. Under taxable muni rates, total ABO liabilities are $3.21 trillion and the gap between liabilities and assets is $1.27 trillion. Under Treasury rates, total ABO liabilities are $5.20 trillion and the gap between liabilities and assets is $3.26 trillion, which is more than 25% of total *federal* debt (roughly estimated at $12 trillion). Note that total outstanding non-pension municipal debt at the state level was $1.00 trillion as of 2008, and 2008 state tax revenues were $781 trillion. If one treats the gap as off-balance-sheet debt, states are actually more than 4 times indebted as their on-balance-sheet accounts reveal.

As discussed extensively in Novy-Marx and Rauh (2009a), the taxable muni yield curve credits each state for its likelihood of defaulting on its liabilities, but probably overestimates the state's ability to do so. However, the Treasury yield curve suffers from the drawback that its yields are very low due to the extreme liquidity and perceived safety of Treasury bonds (see Krishnamurthy and Vissing-Jorgensen

(2008) among other analyses). Partly offsetting this disadvantage is that the COLAs in the benefits imply that yields lower than comparable nominal bonds should be used.

4. 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?

An answer to the question of the optimal level of public debt is very difficult to provide. 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 (2009c) 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 (2009c)).

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 2000s to issue Pension Obligation Bonds (POBs), which raise financing at the tax-preferred municipal bond rate in order to fund state pension funds has lead to IRS rulings deeming these instruments “arbitrage bonds” and disallowing tax deductibility

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 (2009b)). 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 accordingly. Ricardian Equivalence is even less likely to operate, and intergenerational transfers are even larger.

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 the liabilities they have already accrued. 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 would in fact fund all new benefit accruals, It then asks how much additional money would be needed to plug the gap in the underfunding of 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 $100 billion would extend the funds through most of period during which retirees who have already accrued benefits will be receiving payments. Alternatively, states could choose to bring future benefit accruals to zero by freezing benefits at their current levels and re-directing the annual $100 billion in contributions to ameliorating the funding deficit of already-accrued benefits. However, there would presumably need to be some compensating differential to induce public employees to accept such a deal.

It is worth noting that other countries facing pension crises have adjusted payments to all program participants including existing retirees. The US states therefore are in something of a unique situation given the hardness of the claims that fall under the accumulated benefit obligation (ABO): pension payments to those already retired, benefits pledged to former employees who are not yet retired, and benefits already earned under the existing benefit formulas.

Municipal bond investors seem to be aware of these challenges, at least to some extent. The greater the increase in the gap between assets and liabilities during late 2008, the larger the increase in municipal bond spreads, even within ratings categories and controlling for other measures of state fiscal health (Novy-Marx and Rauh (2009c)). What may be less well-understood by taxpayers is the likelihood with which future generations will have to bear the substantial burden of making up pension benefits for previous generations of state employees.

Furthermore, 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.

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.

Lintner, L., 1965, "The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets," *Review of Economic Statistics* 47, 13-37.

Modigliani, Franco, 1961, "Long-run Implications of Alternative Fiscal Policies and the Burden of the National Debt," *Economic Journal* 71, 730-755.

Modigliani, Franco and Merton H. Miller, 1958, "The Cost of Capital, Corporation Finance, and the Theory of Investment," *American Economic Review* 48: 261-297.

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, 2009a, "Public Pension Promises: How Big Are They and What Are they Worth?," University of Chicago Working Paper

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

Novy-Marx, Robert and Joshua Rauh, 2009c, "Fiscal Imbalances and Borrowing Costs: Evidence from State Investment Losses," Kellogg School of Management Working Paper.

Ricardo, David, 1820, *Essay on the Funding System*, Various editions.

Sharpe, W.F., 1964, "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk," *Journal of Finance* 19, 425-442.

Treynor, Jack L., 1961, "Toward a Theory of the Market Value of Risky Assets," Unpublished Manuscript.

U.S. Census Bureau, 2009, "State Government Finances: 2008," Online Data File, http://www2.census.gov/govs/state/08statess.xls.

Figure 1: Benefits Payments Through 2045 Already Owed to Employees

This figure shows projected aggregate cash flows by state governments due to public pension promises, as would be recognized under the Accumulated Benefit Obligation (ABO) method. As in Novy-Marx and Rauh (2009a), 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 and reported discount rate. The distribution of cash flows over years relies on assumptions about wage growth by age, as well as the age-service distribution of workers, based on financial reports from the 10 states with the largest pension liabilities. Benefit formulas, cost of living adjustments (COLAs) and inflation assumptions are derived on a plan-by-plan basis from the CAFRs and the Center for Retirement Research (2006). See Novy-Marx and Rauh (2009a) for details.

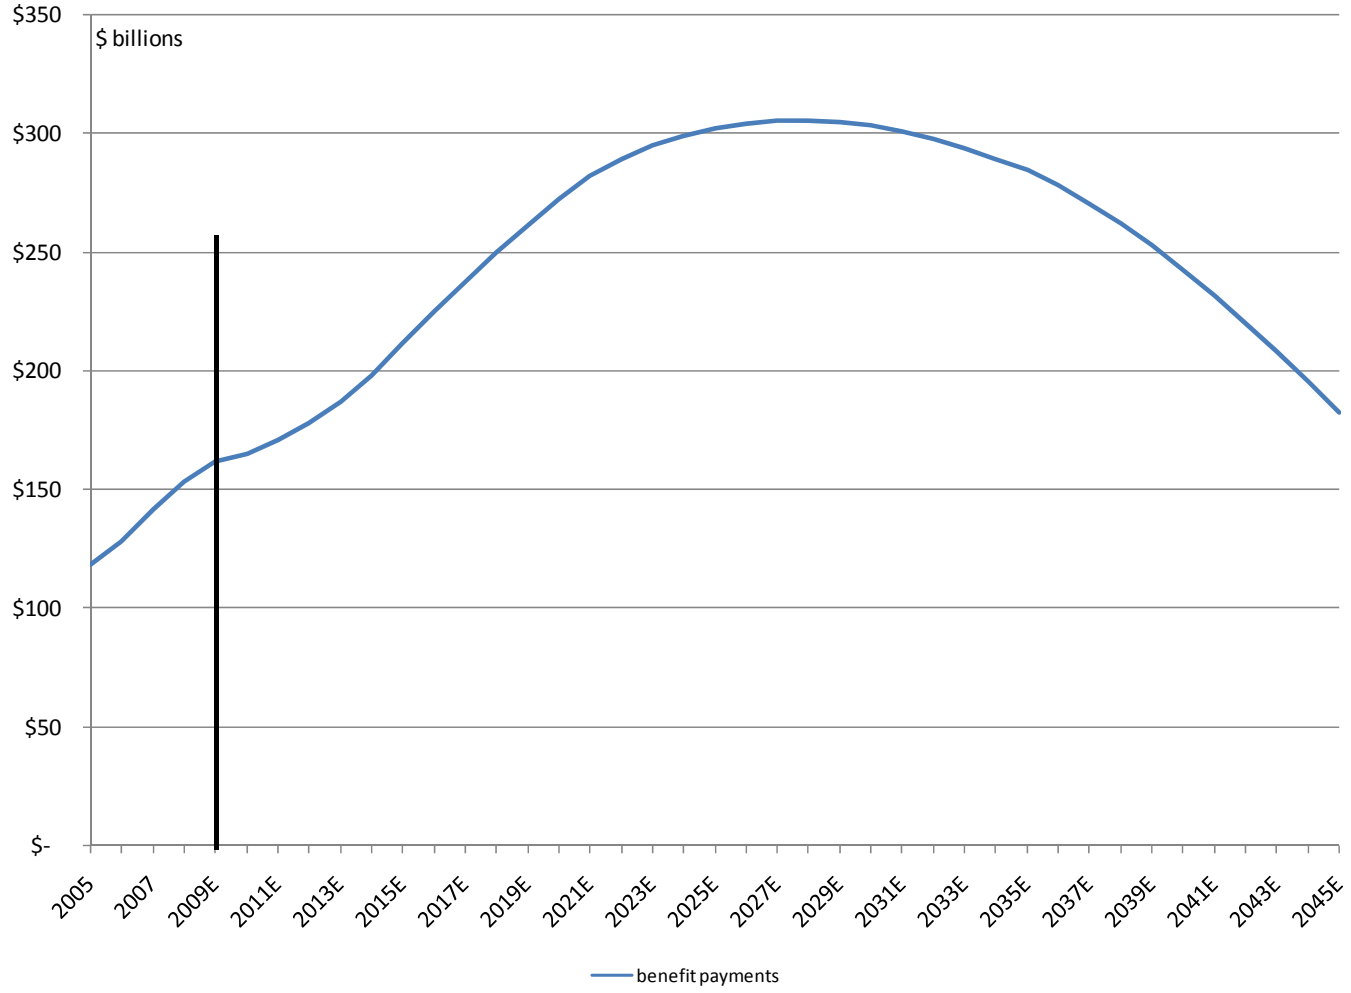

Figure 2: Contributions versus New Benefit Accruals in State Pension Plans

The bars in the figure show aggregate contributions to the 116 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. The line in the graph shows new benefit accruals as measured by the plans' estimated aggregate service cost. This service cost is equal to the change in ABO liabilities at the state-chosen constant discount rate, 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).

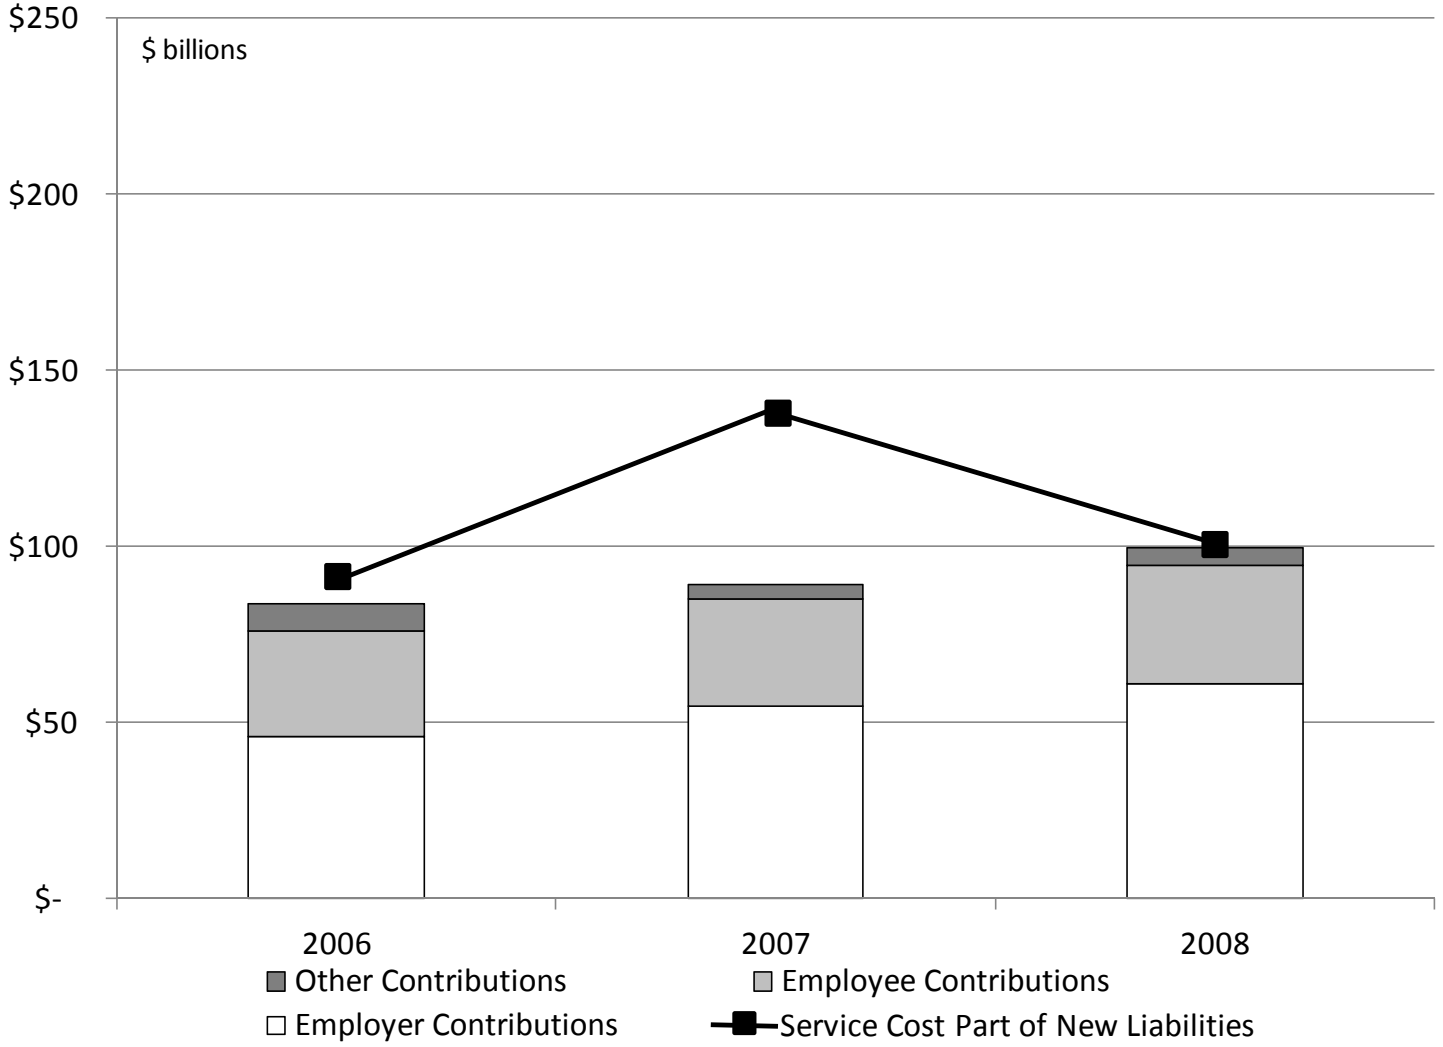

Figure 3: Aggregate State Pension Fund Asset Levels Under Different Scenarios

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 newly accrued liabilities. The middle graph shows assets under the assumption that states contribute an additional $50 billion per year above and beyond newly accrued liabilities during 2010-2020, and the bottom graph shows assets if states contribute an additional $100 billion per year above and beyond newly accrued liabilities during 2010-2020.

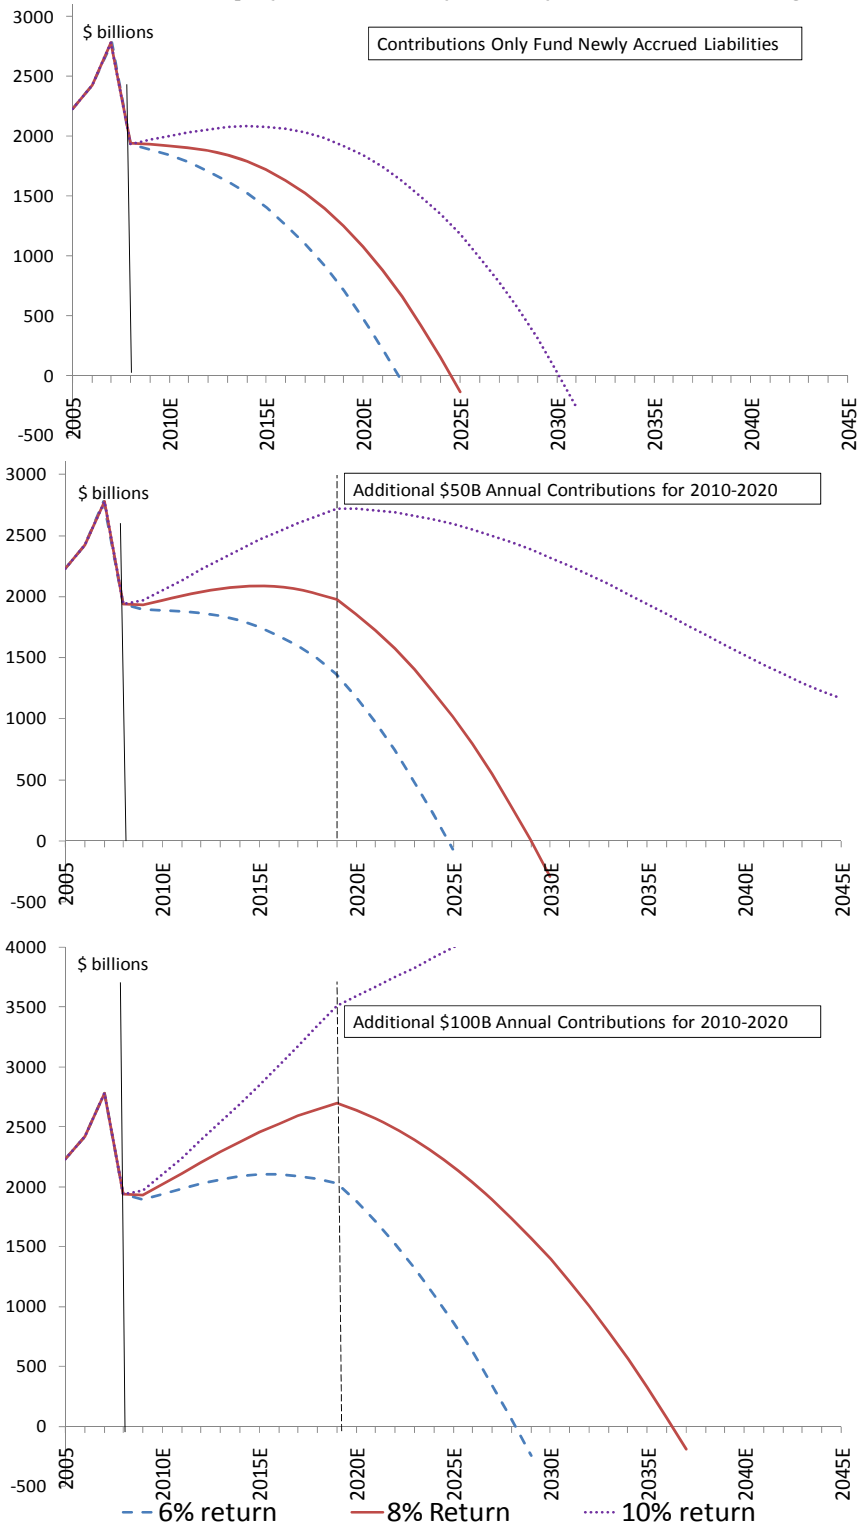

Table 1: Present Value of Aggregate State Public Pension Liabilities

This table summarizes the measures of state public pension liabilities calculated in Novy-Marx and Rauh (2009a). In the left panel, benefits are discounted at state-chosen rates, usually 8%. In the middle panel, benefits are assumed to have equal priority to state general obligation (GO) bonds. They are discounted at municipal bond rates excluding the tax preference, based on the zero-coupon municipal yield curve as of January 30, 2009. In the right panel, Treasury rates are used as an approximation for an appropriate default-free rate, even though Treasury yields likely contain priced inflation and liquidity risks. Under the Treasury discounting procedure, Liabilities are discounted using the zero-coupon Treasury yield curve as of January 30, 2009. In all panels, active worker liabilities are calculated using the Accumulated Benefit Obligation (ABO) accrual methodology.

	Discount Rate		
	State-Chosen	Taxable Muni	Treasury
Total ABO Liabilities	$2.80 trillion	$3.21 trillion	$5.20 trillion
Assets (Estimated at 12/31/2008)	$1.94 trillion	$1.94 trillion	$1.94 trillion
Total ABO Liabilities – Assets	$0.86 trillion	$1.27 trillion	$3.26 trillion

Source: Novy-Marx and Rauh (2009a)