Linking Benefits to Investment Performance in US Public Pension Systems*

Robert Novy-Marx
University of Rochester and NBER

Joshua D. Rauh
Stanford University and NBER

July 1, 2012

Abstract

A large majority of public sector employees in the US receive retirement packages in which by contract they bear no investment risk. Initial retirement benefits depend on salaries and years of service, and cost of living adjustments are either fixed, ad hoc, or linked to inflation. This paper calculates the effect that introducing risk-sharing during either retirement or the working life would have on public sector pension liabilities. Introducing a variable annuity for the retirement phase, modeled on the Wisconsin Retirement System in which positive benefit adjustments are granted only if asset returns surpass 5%, would reduce unfunded accrued liabilities by around 25% while allowing all other plan features to be maintained. It would lower required annual contribution increases for full funding within 30 years by 44%. We then derive alternative policy measures that would have similar effects on costs. These include increasing employee contributions by 11.6% of pay in order to keep benefits unchanged; or giving employees a collective or DC plan with a 9% employer contribution for future service. We discuss these results in the context of models of lifecycle portfolio choice, which suggest that employees should generally prefer to take risk earlier in their lives rather than later.

* Novy-Marx: (585) 275-3914, Robert.Novy-Marx@simon.rochester.edu. Rauh: (650) 723-9898, Rauh_Joshua@gsb.stanford.edu. Rauh gratefully acknowledges funding from the Zell Center for Risk Research at the Kellogg School of Management. We thank Eduard Ponds for helpful comments and David Villa for useful conversations.

1. Introduction

Pension systems in the US are typically either pure defined benefit (DB) plans, in which the employer bears all the investment risk and responsibility for asset allocation, or individual defined contribution (DC) plans, in which risk bearing and asset allocation are the responsibility of the employee. If the goal of a pension system is to provide economic security in old age in a financially sustainable way (Barr and Diamond (2011)), neither of these plan types has succeeded. State and local DB plans in the US have funding arrangements that have created large fiscal liabilities for their sponsors (Novy-Marx and Rauh (2011a)), with unfunded pension debts currently running at over $4 trillion. The shift of public pension funds into alternative assets has also sparked debates as to whether the performance of these asset classes justifies their high fees. Meanwhile, evidence suggests that the 401(k) implementation of the individual DC plan has resulted in employees making suboptimal savings and investment decisions (Brown, Liang, and Weisbenner (2007), Choi, Laibson, and Madrian (2011), Tang, Mitchell, Mottola, and Utkus (2010)) and paying substantial fees for the funds in their individual accounts.

In addressing the problem of unfunded DB liabilities, several states including Colorado, Minnesota, South Dakota, and New Jersey have reduced or eliminated automatic cost of living adjustments (COLAs). In response to underperforming investments, these states have essentially performed *ex post* renegotiations of the pension contract, partially converting real benefit streams into nominal benefit streams. Other states have responded to unfunded liabilities by raising employee contributions, another example of *ex post* cost shifting.

Rather than attempt continual renegotiation of contracts, an alternative is to implement *ex ante* risk sharing. For example, participants in most DB systems in the Netherlands now bear some investment risk. The retirees only receive COLAs if the assets in the fund have performed sufficiently well during both their working lives and during their retirement, an arrangement known as conditional indexation (see Bovenberg and Nijman (2011), Ponds and van Riel (2009)). In some Dutch plans, employees bear even more risk through a collective DC arrangement, where not only the COLA but also the accrued benefit is

a function of asset performance and the sponsor no longer provides a benefit guarantee.[1] The cash balance plans that have replaced a number of traditional DB plans in the corporate sector are a hybrid arrangement analogous to a collective DC plan with a minimum return guarantee during the employee's working life.

Public sector plans in the US rarely involve true risk sharing. In several states, employees receive both DB and DC benefits, and in isolated cases (e.g., Oregon) DC assets are pooled. Nebraska state employees and certain Texas municipal and county employees are in cash balance plans (NASRA (2011)). The Wisconsin Retirement System (WRS) is unique among US public sector DB plans in that post-retirement annuity adjustments explicitly depend on investment returns. Retirees in the Core program of the WRS receive no COLAs, but rather performance-linked benefit increases that are granted only if smoothed asset returns achieve at least a 5% return threshold and all actuarial assumptions are met. In contrast to a collective DC plan with annuitization, in which the participants bear risk during the accumulation phase of the lifecycle, participants in the Core WRS only bear the risk during retirement.[2] We call this feature of the Wisconsin system performance-linked annuity adjustments, or PLAAs.

This paper considers the effects that introducing risk sharing either late in the lifecycle through PLAAs or earlier in the lifecycle through a collective DC arrangement would have on US public pension liabilities, and on the cash flow demands pension systems have on state budgets. Specifically, we begin by measuring the effect that the PLAAs have on benefit cash flows and the present value of liabilities in Wisconsin. We then calculate the effects that indexation of this variety would have if it could be applied prospectively to retiree benefits in the other 49 states, and we derive the employee contribution increases that would have equivalent effects on government budgets. Finally, we calculate how large the employer

[1] Relative to an individual DC plan like a 401(k), however, investment risk in a collective DC plan is pooled both within and across generations of participants in the plan, but is not borne by the sponsor. The United Kingdom Department for Work and Pensions (2008) surveys these arrangements.

[2] As will be discussed in Section 3, the WRS has other option-like features including the participant's option to take a money-purchase benefit based on contributions and investment performance instead of a formula-based benefit. So in fact WRS employees can participate in the upside of investment performance during the accumulation phase, although not the downside. Furthermore, WRS employees have the option to participate in a Variable Fund program that increases the amount of risk borne by the participant during both the accumulation phase and the decumulation phase. The Core benefit represents most of the system assets.

contribution would have to be under the introduction of a collective DC arrangement if similar cost-savings were to be achieved.

Replacing COLAs across the US with PLAAs similar to those in Wisconsin would reduce legacy liabilities by $575 billion and would lower the present value of expected benefit payouts by $912 billion. It would lower the annual required revenue increases to fund state plans within 30 years from $1,375 per household per year to $776 per household per year. Tying benefit adjustments to higher threshold rates of return would of course have stronger cost shifting effects. In the limit, a variable annuity with a high enough threshold and a guaranteed nominal floor is equivalent to eliminating COLAs entirely, which Novy-Marx and Rauh (2011b) show would reduce unfunded liabilities by approximately half. The cost savings to states from moving to PLAAs similar to those found in Wisconsin would be equivalent, on average, to states requiring employees to increase payroll contributions by 11.6 percentage points, i.e. 11.6% of pay.

The PLAA arrangement leaves participants bearing risk only during retirement, not during the time they are working. Standard intuition from the lifecycle portfolio literature (Jagannathan and Kocherlakota (1996), Heaton and Lucas (1997), Viceria (2001), Campbell and Viceira (2002)) suggests that given a choice, individuals prefer to bear risk during the earlier years of their lives instead of the later years. We find that freezing DB plans and implementing collective DC plans in which the employers contribute 9% of pay to the employee's account would achieve the same cost savings as introducing PLAAs for all employees. The collective DC plan could offer either mandatory or voluntary annuitization upon retirement, which if converted at market rates would impose no additional costs on plan sponsors. This analysis assumes that the investment and administrative expenses of running an a collective DC plan would be the same as running the traditional DB plan, which seems reasonable given that in both instances the investments are pooled.

We note that there is substantial heterogeneity across states in these results. States that currently have large COLAs stand to gain the most from implementing variable annuity adjustments. States with relatively high service cost accruals of the DB plan (e.g., benefit factors), those with relatively low current

contribution rates, and those in which employees are already mostly in Social Security benefit relatively more from introducing the collective DC arrangement than from the introduction of variable annuities. While our analysis considers the effects of these reforms as mandatory measures for all workers, there also would be substantial heterogeneity across worker groups in which of the reform options would be preferred. Older participants would certainly prefer employee contribution increases and the introduction of risk-sharing for future service and salary rather than PLAAs. Furthermore, whether participants would actually prefer to bear risk early or late − or whether they would pay to avoid bearing it at all − depends on their risk preferences and the relative number of years they plan to spend working and retired.

We briefly discuss some of the legal issues that states would face in converting COLAs to PLAAs. In states where benefits are protected by "diminished or impaired" language, such reforms (and in fact any reforms to existing pension benefits) would probably not be legal without changes to state constitutions (Monahan (2010)). In others, they may be applicable to some or all participants. Introducing risk-sharing through a collective DC arrangement for future accruals would be allowable as long as the law does not require the state to provide the option to accrue benefits of equal or greater generosity than those previously available to the employee.

To be clear, a move from guaranteed COLAs to PLAAs would achieve cost savings for taxpayers by reducing the present value of expected benefits for existing employees. In the 40% of plans where existing COLAs are linked to consumer price inflation, the reform would leave members exposed to inflation risk that they previously did not bear. In plans where existing COLAs guarantee fixed-rate benefit increases, the relief for taxpayers is achieved by linking the benefit increase to performance, instead of increasing payments irrespective of asset returns. In this latter type of plan, employees currently do not have inflation protection *per se* but simply higher expected levels of benefits. The variable annuities in PLAAs do provide some inflation protection, however, because the performance thresholds are nominal, whereas long run nominal investment performance is correlated with inflation. The collective DC arrangement, on the other hand, would allow for the provision of inflation-indexed annuities, albeit at market prices.

This paper proceeds as follows. Section 2 describes background, literature, and theory on variable annuities and hybrid pension systems including those with conditional indexation and collective DC features. Section 3 discusses the unique features of the Wisconsin Retirement System, and their impact on plan liabilities in Wisconsin. Section 4 calculates the effects that implementing PLAAs in all 50 states would have on the present value of liabilities and on the revenue demands of funding public employee pension promises. In Section 5, we consider the employee contribution increase and collective DC parameters that would have similar budgetary effects. Section 6 concludes.

2. Hybrid Systems: Background, Literature and Theory

2.1 Background on Hybrid Pension Systems and Variable Annuities

Pension systems can be classified along a number of different dimensions. The primary way that pure collective DB and pure individual DC differ is the bearing of investment risk and actuarial (particularly longevity) risk. In traditional DB plans, essentially all this risk is borne by the employer. However, DB sponsors do have an implicit option to increase employee contributions if investments perform poorly or actuarial variables develop unfavorably for the sponsor, so that there is also an element of sharing risk across generations.[3] In individual DC plans, all of these risks are born by the employee, although employees in these plans can generally annuitize (at least in the private market) and therefore choose not to bear longevity risk.

Despite the polarization of the US retirement saving landscape, there is in fact a much wider range of system types and risk-sharing arrangements that could be employed. Table 1 gives examples of pure and hybrid employer-sponsored pension systems according to the distribution of risk on three primary dimensions: investment risk during the accumulation phase, investment risk during the

[3] Additionally, intergenerational risk sharing gives market exposure to younger members of the system that are too wealth constrained to borrow and obtain such exposure directly.

decumulation phase (retirement), and longevity risk.[4] It also shows how in some plan types investment choices are directed by the individual participant, whereas in others they are pooled and made collectively, usually by the employer but with input from employee trustees.

A hybrid plan common among US corporations that previously sponsored traditional DB plans is the cash balance plan. In this arrangement, the employee and employer contributions are pooled and invested by the sponsor, who promises a minimum guaranteed rate of return (Clark and Schieber (2004), Coronado and Copeland (2004)). Investment risk is therefore shared between employer and employee during the accumulation phase. The closer the guaranteed return is to a risk-free rate of return, the more risk the employee bears. If the employee retires and chooses an annuity option, the cash balance is used to purchase an annuity, generally at market rates, so that any investment and longevity risk is shifted to the annuity provider. If the employee retires and chooses a lump sum, then the employee bears both the investment risk and the longevity risk.

A small number of public sector pension systems in the US offer cash balance plans to employees. These include the Nebraska County and State Plan, the Texas Municipal Retirement System (TMRS), and the Texas County and District Retirement System (TCDRS). These plans are described in NASRA (2011). In the Nebraska plan, employees contribute between 4.5% and 5% of pay, with the employer matching approximately 150% and guaranteeing a rate of return on all investments that is at least 5%. The Texas plans work similarly with contribution rates depending on the group, and statutorily-determined guaranteed returns of 5% (TMRS) or 7% (TCDRS) respectively. In examining private sector cash balance plans, Coronado and Copeland (2004) conclude that firms converting to cash balance plans operate in competitive industries with tight labor markets and do not necessarily lead to reduced pension liabilities. Clark and Schieber (2004) find that converting to cash balance plans reduces expected pension payments for longer-tenured workers remaining on the job past age 55, but increases them for workers who quit or are laid off earlier.

[4] Bovenberg (2011) classifies pension schemes overall on two dimensions: state organized / privately organized, and individual choice / mandatory. Our classification does not include state-organized plans that are comprehensively designed to cover all individuals in a society.

The US cash balance plans can be thought of more broadly as a form of collective DC plan with a return guarantee. The aim of collective DC plans is generally to avoid the drawbacks of traditional DB plans and those of individual-accounts DC plans, while including the best features of each. Traditional DB plans are not very portable and reward certain labor market behaviors while punishing others, problems which collective DC plans clearly solve. Traditional DB plans also have service costs that can be difficult or controversial to measure and have generated large unfunded liabilities in both the corporate and public sectors. Collective DC plans can address this issue if they do not contain return guarantees, in which case employees bear the investment risk during the accumulation phase. With return guarantees, the sponsor takes on the downside risk of investment returns below the guarantee. If the plan is not investing in assets that can directly match the guaranteed return then the collective DC plan will still have to measure accruals separately from contributions, and the emergence of unfunded liabilities for the sponsor is still possible. In collective DC plans, the asset allocation is chosen by the sponsor, not by the participant, which avoids problems with poor asset selection by the employee but also requires a one-size-fits-all asset allocation that may not be optimal for all employees.

In addition to the question of guarantees, there are also several ways that collective DC plans can be run. A straightforward arrangement such as a cash balance plan provides a balance for each employee at each point in time based on the contributions made by him or on his behalf and on the annual investment returns. This arrangement does not allow for intergenerational risk sharing, however. A theoretically attractive alternative, which may be more difficult to implement in practice, is to have employees earn shares of a pension trust over their careers, at a price initially set at the time the employee was hired and subsequently adjusted with interest rates. This arrangement provides the employee with more market risk early in the lifecycle, enabling better inter-temporal diversification of investment risk. For a worker with a known retirement date, it is economically equivalent to loaning the employee the current discount value of his future pension contributions (employer and employee), which can be invested today, relaxing the liquidity constraint faced by most young workers. A disadvantage to this setup, however, is that employees can simply quit if the value of the trust falls, reducing the value of their

total compensation below their opportunity cost of providing it. In this paper we focus on the simpler arrangement.

A separate issue from risk sharing in the accumulation phase is risk sharing in retirement, which is closely related to the question of annuitization. There are essentially thee options for annuitization. First, a plan could offer mandatory annuitization, in which the employer provides the annuity. In US public sector plans, this annuity is usually provided directly by the government, but there is also the possibility of the government buying an annuity for the employee at market rates, which would ensure that the employer's liability for retirement is limited only to the amount needed to buy the annuity. Second, a plan could provide the option for voluntary annuitization, in which case the employee chooses between an employer-provided annuity as above and a lump sum. Third, there could be no annuitization at all, in which case the employee takes a lump sum (and could buy an annuity in private markets, although typically does not).

The approach in Wisconsin of providing what we call PLAAs leaves the retirees bearing investment risk, albeit with a floor provided by the system. As will be discussed in the following sections, to the extent that annuities for beneficiaries who are above their floors can be rolled back if others are against theirs, it is in fact other participants that provide the first tier of insurance against benefit payments dropping below their initial levels. Only when all of the positive accumulated adjustments are exhausted are the state and its taxpayers responsible for paying to keep retired beneficiaries at the floor.

The PLAA is in a sense the inverse of a collective DC (or cash balance plan) with annuitization, as far as the timing of the risk-bearing by the different parties is concerned. Under a collective DC (or cash balance) plan with annuitization in retirement, investment risk during the accumulation phase is borne by the participant (or in the cash balance case shared between the sponsor and the participant), whereas during the decumulation phase it is borne by the annuity provider. Under a DB plan with a PLAA arrangement, the investment risk during the accumulation phase is borne entirely by the sponsor and the investment risk during the decumulation phase is borne by the participant.

Risk-sharing through variable annuities in the decumulation phase of retirement savings is treated in the literature on variable annuities. Brown, Mitchell, and Poterba (2001) use a life-cycle model to show that many consumers would find it welfare-enhancing to hold a portion of their retirement portfolio in an equity-linked annuity product. Similarly, Horneff, Maurer, Mitchell, and Stamos (2009) find that variable annuities are part of an optimal dynamic portfolio choice solution when an investor seeks to benefit from holding both equity and longevity insurance. Brown and Poterba (2006), however, emphasize that historically the bulk of variable annuity products available to retail investors are not in fact converted into life annuities, and that instead these products are viewed to some extent as tax-preferred mutual funds. The PLAA setup we examine is one type of variable annuity that could be offered to beneficiaries as part of a pooled DB plan.

Finally, given the choice between bearing investment risk in retirement and during their working lives, there are reasons why workers may prefer the former. Certainly the intuition of standard life-cycle models (Jagannathan and Kocherlakota (1996), Heaton and Lucas (1997), Viceria (2001), Campbell and Viceira (2002)), in which young workers' total wealth is dominated by bond-like future labor income, suggests that workers should prefer market exposure when they are young to when they are old. Moreover, the ability to delay retirement provides a potential hedge against investment losses incurred over the working loss, something that has been observed empirically (Gustafson (2012)), further arguing for early exposure. Even so, workers may still prefer late market exposure. The primary reason for this is that post retirement exposures typically expose the worker to less wealth uncertainty, even for the same level of market participation. Most people simply work for a longer period than they spend in retirement, including in the public sector (excepting public safety workers).[5] Moreover, late exposure only accrues over time, and is further mitigated by the very act of decumulation. The initial annuity payments a retiree receives have almost no exposure to investment risk, and it is only late in retirement that investment performance typically has a substantial impact on the level of the retirement benefits.

[5] And in any case, a system where individuals have a longer retirement than period of work would seem difficult to sustain financially.

2.2 Valuing the Option Features of Hybrid Plans

In this section we develop formulas and a methodology for valuing the option features of hybrid plans, focusing on the PLAA arrangement. One way to conceptualize risk-sharing during the decumulation phase is to assume that the first benefit payment in retirement is fixed by formula but that each successive benefit payment will be determined in part by the evolution of investment returns and/or actuarial experiences. The cost of any given risk-sharing plan is given by capitalizing the initial benefit using an appropriate annuity factor.

Consider as a very simple example an employee whose first-year benefit based on his service and salary would be $30,000, and for whom a 25-year benefit will be provided (we discuss the lifetime annuities below). Assume further that the real yield curve is flat, with TIPS of all maturities yielding real returns of R_{real}. The cost of providing a benefit indexed to the CPI, ignoring mortality risk, would simply be $30,000 times the annuity factor $(1/r)*(1-(1/(1+r)^{25}))$, evaluated setting $r = R_{real}$. It is clear that this is the cost of providing this benefit, since the sponsor could implement the provision of the annuity by investing in TIPS worth $30,000 times this annuity factor, and the participant's expected annuity adjustment is expected inflation, $E[\pi]$.

Alternatively, the plan could promise a benefit that would rise if investment performance were above a threshold (say 5%) and fall if it were below the threshold. The cost of providing this benefit would be $30,000 times the annuity factor $(1/r)*(1-(1/(1+r)^{25}))$ evaluated at $r = 5\%$. This can be seen as the cost of providing the benefit by imagining that the sponsor transfers this amount to a separate account for retirees. Each year the benefit adjusts, such that the remaining payments capitalized at a 5% hurdle rate yields the remaining account balance, or alternatively growing the payment by the realized return minus 5%. The higher the threshold, the less valuable this payment would be. Because payments are made out of the account at the threshold rate, higher thresholds reduce the capital used to fund future payments, retarding their growth rate. In effect, one can think of the threshold as defining a capitalization

rate (r) that determines how much money will be transferred to the separate retiree account, and the transferred amount is the present value of what the retiree receives.[6]

This is one type of variable annuity. Of course, a performance-linked annuity only protects the beneficiary against inflation to the extent that the performance of the assets in the pension fund is correlated with inflation. Under this arrangement, the participant receives an expected annual annuity adjustment equal to the difference between the expected asset returns and the capitalization rate: $E[R_{assets}] - r$. Since $E[R_{assets}] = R_{nominal} + \lambda$, where λ is the expected excess return on the assets, and $R_{nominal} = R_{real} + E[\pi]$, where π is the inflation rate, that implies that the expected COLA is expected inflation minus the difference between the capitalization rate and the expected real return on the assets: $E[\pi] - (r - (R_{real} + \lambda))$, as opposed to $E[\pi]$ under full inflation indexation. The expected performance-linked annuity adjustment thus exceeds the expected adjustment under full indexation whenever the expected real return on plan assets exceeds the capitalization rate.

The annual annuity adjustment the participant expects to receive is lower, however, under the risk-neutral measure. Under this measure the expected adjustment equals the difference between the risk-adjusted expected asset returns (i.e., the nominal risk-free rate) and the capitalization rate: $R_{nominal} - r$. The risk-neutral expected COLA is thus expected inflation minus the difference between the capitalization rate and the real rate of interest: $E[\pi] - (r - R_{real})$. So the performance-linked annuity adjustment is less valuable than an inflation-linked annuity, i.e., less costly to provide, provided the capitalization rate exceeds the real rate, $r > R_{real}$.

Having begun with these simple examples, we now introduce a term structure of real interest rates and lifetime annuities with survival probabilities. The fully inflation-indexed annuity can be defeased by buying a portfolio of TIPS today. For a cohort of retirees of a given age and average benefit payment the plan needs to hold TIPS strips maturing each year in the future in proportion to both the initial payment

[6] The capitalization calculation above reveals that inflation-linked payments can also be implemented as a performance-linked annuity with a threshold rate equal to the real interest rate and the assets invested in TIPS.

and the proportion of the retirees expected to survive until the payment date. The cost of providing these benefits are only as much as the assets that are set aside. The portfolio

$$b_i \sum_{T=0}^{\infty} S_{a,a+T} \tilde{B}_{0,T},$$

where b_i is the initial benefit payment, $S_{a,a+T}$ is the probability of surviving from the retiree's current age of a to $a+T$, and $\tilde{B}_{0,T}$ is the cost of an inflation linked zero-coupon bond, i.e., a TIPS strip maturing in T years that in T years pays \$1 plus the CPI adjustments realized between today and the maturity date.

One source of risk in the above examples is mortality risk. Individual mortality risk will not matter if the mortality hazard assumptions made by the sponsor are correct on average, but there is a risk that the assumptions about the distribution of mortality will overall be inconsistent with experience. In the standard DB arrangement and in the above formulation, the sponsor bears the full mortality risk; whenever a typical state or local government sees its retirees living longer than expected, the taxpayers are still obligated to pay the full pensions.

We detail an alternative to this below. Here, we note that if there is no annuity adjustment to reflect differences between assumed and realized mortality, the above defeasement strategies could fall short due to mortality risk. That is, if individuals overall live longer than projected, the amount of money set aside to fund retiree annuities would be inadequate unless the annuity is adjusted to reflect that.

The PLAA with a hurdle rate of r makes benefit payments each year T in the future of $b_i S_{a,a+T}(1-r)^T R_{0,T}^{assets}$, where $R_{0,T}^{assets}$ is the gross realized T-year return on plan assets. These cash flows can be replicated (or defeased) with the portfolio that holds

$$b_i \sum_{T=0}^{\infty} S_{a,a+T}(1-r)^T$$

dollars of the plan's assets.

In the example above, there is no floor on the benefit. The benefit can fall as low as needed in order to keep the resources in the fund adequate to pay the annuities. Systems can also build in benefit floors, which give retirees some protection from returns below the threshold, with a guarantee by the

sponsor or by participants whose annuities have risen above the benefit floor. One simple floor would be to promise that retirees will receive no less than the benefit given by their benefit formula, and that the sponsor would guarantee this floor on an individual-by-individual basis. In this case the plan makes benefit payments each year T in the future of

$$b_i S_{a,a+T} \max\left(1, (1-r)^T R_{0,T}^{assets}\right).$$

These payments can be replicated with the portfolio

$$b_i \left(\sum_{T=0}^{\infty} S_{a,a+T} K_T^{-1} \right) V_0 + b_i \sum_{T=0}^{\infty} S_{a,a+T} K_T^{-1} P(K_T, T) ,$$

where V_0 is the plan's assets and $P(K_T, T)$ is a T-year European put option on the plan assets struck at $K_T = V_0/(1-r)^T$. The value of the floor, i.e., the put portfolio, is increasing in the capitalization rate and more valuable when interest rates are low.[7]

The above discussions assume that both aggregate mortality risk and the investment risk inherent in providing the floor are borne by the sponsor, in this case state taxpayers. However, an alternative arrangement could be that participants who have received performance-linked increases in the past provide either the mortality insurance, or the investment return insurance for those against the floor, or both. The sponsor will then not be liable for the guarantee, or for incorrect assumptions about mortality, until all of the retirees' benefits have hit the floor.

Therefore, retirees with accumulated performance-linked benefit increases could provide co-insurance so that other retirees do not fall below the benefit floor in the case of poor asset performance. And these retirees with surpluses could also provide co-insurance on mortality. If retirees live longer, then those with accumulated performance-linked benefit increases will see their benefits reduced to pay for that longevity — at least until their benefits reach the floor, at which point it becomes the sponsor's responsibility.

[7] The first fact follows because the payoff to $K_T^{-1} P(K_T, T)$ is increasing in the strike, K_T, which is itself increasing in the capitalization rate, while the second holds because put options have negative rhos.

Due to the complexities of calculating cost when members of the pool co-insure one another, our methodology assumes that the sponsor provides guarantees on an individual-by-individual basis. That is, for both asset returns and mortality, the sponsor bears the full burden of providing the floor of the individual's formula benefit.

As we discuss below, the co-insurance approach is in fact implemented in Wisconsin. Ignoring the approach overestimates the cost of providing benefits if the co-insurance is implemented.

The variable annuity approach differs from true conditional indexation, in which the annuity is asset linked but cannot go below a floor of the nominal benefit, nor above a ceiling of the accumulated CPI increases. True conditional indexation can be implemented with a bond portfolio worth

$$b_i\left(\sum_{T=0}^{\infty} S_{a,a+T}\, K_T^{-1}\right)V_0 + b_i\sum_{T=0}^{\infty} S_{a,a+T}\, K_T^{-1} P(K_T,T) - b_i\sum_{T=0}^{\infty} S_{a,a+T}\, K_T^{-1} X\left(V_0, K_T\tilde{B}_{0,T}, T\right)$$

where $X\left(V_0, K_T\tilde{B}_{0,T}, T\right)$ is an exchange option giving the holder the right at maturity to trade the plan asset, currently worth V_0, for zero coupon inflation linked bonds with a face equal to K_T (today's dollars), currently trading at $K_T\tilde{B}_{0,T}$. These exchange options are priced using Black's formula for options on futures, where the underlying is the future price of the receivable, the strike is the forward price of the deliverable, and the relevant volatility is that of the ratio of these two forward prices.

Calculations of the value of bond portfolios and options that replicate the hybrid plan features are done as described in this section. Options prices are calculated using Black-Scholes with a 12% volatility assumption.

2.3 Methodology for Revenue Demands

In this paper we are relying on similar procedures to Novy-Marx and Rauh (2011a) in determining the cash flow benefit payments that states will have to make before variable annuities are implemented. One difference is in the treatment of inflation. We consider real cash flows, deflating nominal cash flows forecast from Novy-Marx and Rauh (2011a, 2011b) using the inflation assumption built into the forecast nominal benefit payments. To discount liabilities, we use real Treasury yields (based on TIPS) as of December 31, 2010 on the deflated cash flows. In particular, we forecast real

liability cash flows using the uniform inflation assumption of 2% per year, adjusting COLAs and wage growth assumptions appropriately to reflect the differences between this rate and the plans' stated inflation rate assumptions. This inflation assumption, based on today's market-implied inflation expectations, is substantially below that used by the states themselves. Implementing it reduces the present value of benefits by around 15%.[8]

There are three liability concepts that we consider. The first measure, known as the Accumulated Benefit Obligation (ABO) focuses only on payments that have already been promised and accrued. In other words, even if the pension promises were frozen at today's levels of service and salary, states would still contractually owe these benefits. The ABO liability is not affected by uncertainty about future wages and service, as the cash flows associated with the ABO are based completely on information known today: plan benefit formulas, current salaries, and current years of service. However, the ABO liability assumes that the COLAs are preserved.[9]

If workers receive their marginal product in total compensation (wages plus pension benefits), the ABO is the only relevant concept since it measures the benefits that employees have actually earned (Bulow (1982), Brown and Wilcox (2009)). The ABO is a "narrow" measure in that it does not recognize any future wage increases or future service that employees are expected to provide, even though such wage increases and service are to some extent predictable. Moreover, the ABO obligation is independent of wage risk, which simplifies the valuation.

The two broader measures (EAN and PVB) account to different extents for the fact that benefits will continue to accrue due to the future salary and/or service of existing workers. They assume that the pension system will not be frozen today and aim to reflect some portion of actual expected benefits.

The broadest measure, the PVB, represents a discounted present value of the full projection of the cash flows that actuaries expect the state to owe. The PVB method does not credit the state for the fact

[8] State actuaries may well be making other assumptions that offset this and bias the liability downward. For example, many use outdated mortality tables that do not accurately reflect today's expected lifespans.
[9] As such one element of uncertainty in ABO cash flows is the realization of future inflation, at least in systems where the COLAs are linked to CPI inflation.

that it might have some ability to limit benefit accruals. The EAN recognizes a fraction of the PVB. Mathematically, the EAN method recognizes the PVB in proportion to discounted wages earned to date relative to discounted expected lifetime wages. In practice, this procedure accounts for some portion of future benefit accruals due to both wages and future service. See Novy-Marx and Rauh (2011a) for further discussion.

In addition to estimating the effects of introducing variable annuities on the present value of benefits, we also study their effects on the revenue demands required for full funding of state and local benefits in 30 years. These calculations draw on the methodology in Novy-Marx and Rauh (2011c). In this analysis, we assume that the real value of assets grows at the point on the TIPS yield curve that corresponds to the average duration of real liabilities (21 years), which is 1.7%. This assumption implies that the nominal value of assets grows at inflation plus 1.7%. Again, we use real Treasury yields (based on TIPS) to discount cash flows to calculate the change in the present value ABO liability resulting from an additional year of work.

Specifically, to calculate these revenue demands, each year plan assets are assumed to grow at a real rate of 1.71%, the 21 year zero-coupon TIPS yield as of 31 December 2010, where this maturity is picked to match the duration of the real pension liabilities at the corresponding yield. This is the real rate that may safely be achieved when assets are picked to match liabilities, and is equivalent to assuming that assets will grow at inflation plus 1.71%. Assets are then reduced by the benefit payments made that year, to reflect outflows to plan participants.

To these assets we add the contributions from plan participants, which are assumed to be a constant fraction of wages. For each state the contribution rate for plan participants is taken from the data, and averages just under 6%, though there is a great deal of variation across states. Plan participants' aggregate salaries are taken from the model, and account for mortality, retirement, and wage growth.

Finally, we add the contributions from employers, less the cost of new service accruals. State and local governments are assumed to contribute a constant fraction of total adjusted payrolls for the next thirty years, the "amortization rate." Total payrolls, as well as GSPs, are assumed to grow at a constant

real rate, and we consider several different scenarios: growth consistent with individual states' experiences over the last ten years, growth consistent with the national experience over the last ten years, and each of these scenarios reduced by one percent.

Total assets *T+1* years in the future, A_{T+1}, are therefore given by

$$1.0171A_T + \left(AR^* - \left(c^{NC} - c^{employee}\right)\right)(1+g)^T W_0^{total} - B_T\,,$$

where *AR** is the amortization rate (our primary object of interest), c^{NC} is the normal cost rate (service cost relative to wages), $c^{employee}$ is the employee contribution rate, g is the assumed growth rate in the state's economy and government sector, W_0^{total} is total wages today, and B_T is the deflated time-T benefit cash flows to retirees currently recognized under the accounting methodology (ABO or EAN). We search for the amortization rate *AR** such that assets thirty years in the future are just sufficient to pay the remaining recognized benefit payments owed to current workers, i.e., such that

$$A_{30} = \sum_{t=0}^{\infty} \frac{B_{30+t}}{(1+r)^{t+1}},$$

where *r* is picked to match the 21 year TIPS rate of 1.71%. If the assets together with expected investment earnings are insufficient to pay remaining future benefit obligations, then the algorithm tries a higher employer contribution over the next thirty years. If they are more than sufficient, then we try a lower rate. The algorithm searches until it finds the rate that just fully amortizes the legacy liabilities over the thirty year period.

We also consider the possibility of introducing a new collective DC plan for future accruals, in which all benefit accruals are stopped, including for current workers. The cost effects of this for the government are essentially the effects of a hard freeze of DB benefits. No accumulated benefits are taken away, but employees stop accruing defined benefits with additional years of service and salary increases. Furthermore, the administrative costs of managing the collective DC plan and the traditional DB plan are likely to be very similar, since the collective DC plan does not involve the creation of individual accounts or any new layers of fees. In our modeling of the introduction of a collective DC plan, we assume that the

governments need only pay off today's unfunded ABO liability over 30 years. All employee contributions that were previously going to DB plans would go to DC plans, and employers would contribute to the collective DC plans at a specified rate.

To model the introduction of the collective DC plan, the pension cost of future work is assumed to be that of a DC plan with the specified employer contribution, plus the full cost of providing Social Security to new workers in those systems that do not currently enroll workers in Social Security. The cost of Social Security is 12.4% of payroll, which generally is split equally between employers and employees, but our analysis is based on the notion that workers not in Social Security would require pay increases of 6.2% to pay their share, so that the cost of both the employer and employee share would effectively be paid by the employer. This is similar to the hard freeze modeled in Novy-Marx and Rauh (2011c).

The amortization rate under the introduction of the collective DC plan is calculated by modeling total assets *T+1* years in the future, A_{T+1}, as given by

$$1.0171A_T + \left(AR^* - c^{employer(DC)}\right)(1+g)^T W_0^{total} - B_T^{ABO}.$$

We again search for the amortization rate *AR** such that assets thirty years in the future are just sufficient to pay the remaining benefit payments owed to participants of the old, frozen DB plans, i.e., such that

$$A_{30} = \sum_{t=0}^{\infty} \frac{B_{30+t}^{ABO}}{(1+r)^{t+1}}.$$

Readers are referred to Novy-Marx and Rauh (2011c) for further details on the mechanics of revenue demand calculations.

3. Impact of Wisconsin Features on Wisconsin's Liabilities and Benefit Cash Flows

3.1 Features of the Wisconsin Retirement System (WRS)

In this section we explain the features of the WRS. This discussion draws on documents by Wisconsin Department of Employee Trust Funds (2010), the Wisconsin Legislative Audit Bureau (2010), and State of Wisconsin Investment Board (2010).

The WRS has two sub-programs, the Core Fund program and the Variable Fund program. Both entail PLAAs in retirement, although the Variable Fund has a more aggressive investment strategy and therefore annuity benefits are more volatile. Participants in the Variable Fund also share risk with the sponsor during the accumulation phase. Members who wish to participate in the Variable Fund have a one-time option to elect that 50% of their contributions go into the Variable Fund. As of December 2009, 19.3% of WRS employees had made this election. Variable Fund assets as of December 2009 were $5.0 billion and Core Fund assets $67.8 billion. We focus on the Core Fund features in this paper, as the Variable Fund is still relatively small.

At retirement, participants in WRS initially receive the higher of either a *formula benefit* or a *money purchase* benefit. The formula benefit is a standard DB benefit calculated as [benefit factor] x [final average monthly earnings] x [years of service] x [age reduction factor for early retirement]. For service since 1999, the benefit factor has been 1.6% for general employees and teachers, 2.0% for public safety employees covered by Social Security, and 2.5% for public safety employees not covered by Social Security. Table 2 shows these parameters as well as normal retirement ages, normal retirement levels of service, and age reduction factors for early retirement.

The money purchase benefit is an annuity amount that WRS offers to provide based on the balance of the employees' contributions grown by the investment returns that have been realized during the accumulation phase. If the assets in the fund have performed particularly well, the initial money purchase may be higher than the initial formula benefit. Once the initial benefit has been set by either the formula or the money purchase calculation, it remains the basis of the participant's annuity for life. The benefit can rise through PLAAs but it cannot fall below this initially set level.

The money purchase option is valuable to the participant when high asset returns are achieved, and limits the ability of WRS to subsidize the pensions of participants who have contributed during an era of poor investment performance with the performance of those who have contributed during an era of strong investment performance. This is an important aspect of WRS, although the focus of this paper is the risk-sharing features during the decumulation phase. It is important to emphasize that within WRS,

cost savings to taxpayers from risk-sharing with public employees in retirement may be in part offset by the costs of allowing the money purchase option at retirement.[10] When we consider the cost burden reduction effects of implementing PLAAs in other states, we assume that the states do not simultaneously introduce this money purchase option for the setting of the initial benefit.

Once the initial benefit has been set, an amount is transferred from the active account to the retiree account that pays for a performance-linked annuity starting with this initial benefit under actuarially estimated mortality assumptions at a 5% investment return threshold for benefit increases. This transfer is the basis of the claim often made about the WRS that the retiree portion is fully funded.[11]

There are, however, two caveats to this claim. First, the amount transferred does not include the value of the guaranteed floor of the initial benefit. If other retirees have amassed substantial PLAAs, then the additional liability of the state is very small. If most retirees are at the floor, then the additional liability could be quite large. Second, the transfer to the retirement account may involve a transfer out of the account for active members that is greater than employee contributions grown by the investment return. When an employee reaches retirement, if he chooses the money balance option then the sum total of his contributions grown at the realized investment rate of return is transferred. If he chooses the formula benefit, it is by revealed preference due to the fact that the formula benefit is worth more than the contributions grown at the realized investment return; in order to fully fund the retiree portion, it is therefore necessary to withdraw more funds from the active account than have been received by invested employee contributions.

3.2 Effect of Hybrid Features on the Liabilities of the Wisconsin Retirement System (WRS)

Table 3 documents the effect that having PLAAs has on the liabilities of the Wisconsin Retirement System. If instead of the variable annuity the system provided a guaranteed linked COLA to the CPI, the ABO liability would be $106 billion and the EAN liability would be $119 billion. If WRS

[10] We provide a rough calculation of this in Section 4.

[11] For example, David Villa, Chief Investment Officer of and Chairman of Investment Committee, states this in an interview: "After retirement, investment risk is shifted to the retirees. Due to this adjustment process, the retiree portion — which currently makes up half of the Wisconsin retirement system — is, by definition, always fully funded." http://www.kellogg.northwestern.edu/kwo/sum11/features/index.htm

had a PLAA with no floor, the ABO liability would be $81 billion and the EAN liability would be $90 billion. Accounting for the floor raises the ABO liability to $94 billion and the EAN liability to $106 billion. None of these liability calculations account for the money-purchase option, however.

It is important to note that these are liabilities from the taxpayers' perspective. That is, as far as taxpayers are concerned, the variable annuity feature reduces the $106 billion ABO liability to $94 billion. This only occurs because workers in the WRS, unlike in any other major US public pension system, share the risk of poor investment performance. Unlike in any other major US public pension system, their annuities can and do decline when the fund investments perform poorly.

The fact that the system transfers assets from the active account to the retiree account upon the retirement of the employee, in conjunction with the retirees bearing most of the investment and mortality risk, means that the unfunded liability represented by retired workers is due solely to the states' backstop of benefit payments at their initial levels. The cash flows from the active account come from future transfers to the retirement account as workers retire, and these transfers have a significantly shorter duration than the benefit payments from traditional defined benefit pension plans. This reduces the sensitivity of liabilities to the discount rate, and for a given level of real economic liabilities increases the amount that actuaries require to be set aside in a given year.[12]

4. Applying PLAAs to State Pension Systems Across the US

4.1 Legal Issues

Legal restrictions on changes to public sector employee pension benefits are examined in Brown and Wilcox (2009) and Monahan (2010). Several states, among them New York and Illinois, specify that pension rights are set when employment begins and may not be diminished or impaired. This approach goes beyond federal rules for private employers, which clearly protect accumulated benefits only,

[12] We note that the assumptions in WRS about worker age, wage growth and retirement behavior suggest that WRS liabilities are of shorter duration than those of other states, independent of the variable annuity system, primarily because Wisconsin employs extremely modest wage growth assumptions, which are often insufficient to keep up with expected inflation.

although those accumulated benefits are generally construed to include the right to promised COLAs. While the exact interpretation of these clauses have only been tested in courts to a relatively limited extent, it seems likely that a move from COLAs to PLAAs would face successful legal challenges in states such as New York and Illinois on the grounds that benefits had been diminished or impaired.

As reviewed by Monahan (2010), other states either attempt to protect accrued benefits or provide some other kind of protection such as specifying that pension formulas constitute contracts, changes to which must be justified by "an important public purpose". As a result, in many states the extent of protection is unclear. Two states (Indiana and Delaware) use language that describes the pension as a gratuity (Brown and Wilcox (2009)), which would seem to permit a variety of changes.

New Jersey, Colorado, Minnesota, and South Dakota have all implemented changes to COLAs in recent years, and suits against these provisions have not proceeded. Overall, it seems likely that moving from COLAs to PLAAs would survive legal challenges in many but not all states.

4.2 Effects on the Present Value of Liabilities

Table 4 shows the liability effects of applying PLAAs in all 50 states under a variety of different parameters and measures. These figures are only for state-sponsored systems, not locally-sponsored systems, and are estimated as of December 2010. To bring estimated liabilities to December 2010, we begin with the database through June 2009 used in Novy-Marx and Rauh (2011a). We observe that stated liabilities grew at a 5.52% annual rate between plan years 2007 and 2008, and at a 5.51% annual rate between plan years 2008 and 2009. Given the stability of this growth rate, we applied a 5.5% annualized growth rate to liabilities between June 2009 and December 2010, in order to predict the value of what stated liabilities under the systems' own accounting methods would be if they were disclosed as of December 2010.

The results in the first column of Table 4 show the status quo ABO at $4.63 trillion, the EAN at $5.10 trillion, and the PVB at $6.77 trillion. These figures are very close to what is presented in Novy-Marx and Rauh (2011a). As explained in Section 2.3, relative to that paper we assume that liabilities have

continued to grow at a 5.5% annualized rate through 2010, and also implement the lower inflation assumption.

The top panel Table 4 shows how the effects of implementing PLAAs without a floor to active workers only (left panel) and to all participants including retirees (right panel). The bottom panel shows analogous calculations for implementing PLAAs with a nominal floor equal to the initial formula benefit. Obviously converting from COLAs to variable annuities has a greater cost reduction effect for taxpayers if it can be applied to all participants including retirees than if it must be restricted to active employees only. Higher hurdle rates also lower liabilities, as positive benefit adjustments become less likely and negative adjustments more likely. Note that when these hurdle rates are applied to all participants including retirees, our calculations assume that those retirees who are currently receiving benefits above their nominal formula benefit will not see their benefits decline below the current level.[13]

Implementing a 5% hurdle for actives only with a floor reduces the ABO liability from $4.63 trillion to $4.33 trillion. Applying this to all participants reduces it to $4.05 trillion. If plans were changed in this way, it would reduce ABO liabilities by $575 billion, and PVB liabilities by over $912 billion. Based on state asset values of $2.4 trillion, unfunded ABO liabilities would fall by around 25%. If a 7% hurdle were implemented with a floor, it would reduce ABO liabilities by $778 billion, and PVB liabilities by $1.21 trillion. Unfunded ABO liabilities would fall by around 35%.

4.3 Effects on the Per-Household Revenue Demands Required to Achieve Full Funding

Table 5 shows the amount of increased revenue per US household needed to achieve full funding of state pension systems in 30 years on an ABO basis, both under the status quo and if variable annuities are introduced.

The Status Quo lines shows these revenue demands under the status quo, as in Novy-Marx and Rauh (2011c), although here we present figures for state-sponsored plans only. For all states combined, the status quo revenue demand is $1,375 per household per year. That is, to achieve fully funded state-sponsored plans in 30 years, the increased contribution per US household would have to rise by $1,375

[13] That is, the newly implemented floor would be set at the current annuity payment, not at the initial benefit.

per household per year starting today and grow with the nominal size of the public sector. For the five states with the highest revenue demands (New Jersey, Oregon, Rhode Island, Ohio, and Wyoming), this required increase is only $2,154 per household on average. For the five states with the lowest revenue demands (West Virginia, Arkansas, Arizona, Utah, and Indiana), the required increase is only $549 per household on average.

The upper panel of PLAA results shows the effects of introducing PLAAs without any floor under benefits, and the lower panel shows the effects of introducing PLAAs with a nominal floor of the initial formula benefit. The scenario in the box in the lower panel is the one closest to the Wisconsin Retirement System's design, in which there is a 5% investment return threshold and a nominal floor equal to the initial formula benefit.

Table 5 shows that there is rather substantial cost shifting from implementing the variable annuities, particularly for all participants, and even more so if there is no floor provided. The WRS scenario in the box results in revenue demands declining from $1,375 per household on average across all states to $776 per household on average across all states. This is a 44% decrease in required contribution increases for full funding. For the five states with the highest revenue demands to begin with, the decline is 50%.

Implementing higher thresholds would have even greater savings for the state but makes it even less likely that beneficiaries will see increases. Those who are currently receiving benefits above the nominal benefit due to accumulated COLAs would be likely to see decreases in their benefits. In the cases of the states with the lowest revenue demands to begin with, implementing variable annuities for all participants with sufficiently high thresholds and no floor would even allow state governments to reduce their contributions over current levels and still achieve full funding in 30 years.

There is substantial heterogeneity in the extent to which implementing variable annuities would reduce the burden on taxpayers. Implementing PLAAs has the greatest effect in states where the pre-existing COLAs are high.

5. Comparison to Other Reforms

5.1 Employee Contribution Increases

Table 6 performs a similar exercise to Table 5 but focuses on revenue demands as a percent of public sector employee payroll. This analysis allows for a comparison of the potential cost effects of implementing variable annuities with the potential cost effects of increasing the percentage of payroll public sector employees must contribute, a popular reform proposal.

The first row of Table 6 shows that under the status quo, contributions by public sector employees would have to rise by 26.5% of their payroll to put state systems on a path to full funding within 30 years, consistent with Novy-Marx and Rauh (2011c). For the five states with the highest increased contribution requirement on this scaling (Colorado, Pennsylvania, Michigan, Illinois, and Oregon), it would require a 39.9% additional payroll tax on public sector employees, and for the five states with the smallest requirement it would be 12.3%.

The panel called PLAAs without floor shows the amount by which these figures improve if PLAAs could be implemented without any benefit floors provided by the government. This results in quite substantial cost shifting to participants, even more so if the adjustments were to be implemented for all participants as opposed to actives only. Higher hurdle rates also have a greater cost-shifting effect. For example, a 5% threshold for all participants with no floor reduces the required increase by 14.9% of payroll for all states, and by 24.3% of payroll for the five with the highest starting requirement. Over half of the unfunded liability would therefore be removed for taxpayers, although beneficiaries would be exposed to significant risk that benefits could be cut below the initial formula benefit.

Even if a nominal floor of the initial formula benefit is implemented, however, the introduction of the variable annuity with 5% threshold for all participants would still reduce the required contribution increases by 40-50%, from 26.5% of payroll to 15.0% of payroll for all states, and from 39.9% of payroll to 19.8% of payroll for those with the highest starting requirement. The implementation of this sort of variable annuity is therefore equivalent in savings to taxpayers of requiring public employees to pay 11.6% of payroll more towards their pensions.

Introducing a WRS-style PLAA and increasing employee contribution rates by 11.6% of pay would therefore have the same effect on government budgets. However, we note that offering this choice to *each employee* would obviously have much smaller budgetary impact, as those closest to retirement would chose employee contribution increases and those far away would be more likely to prefer PLAAs. The budgetary impact calculation assumes that the plan (either PLAAs or employee contribution increases) is adopted for all employees.

5.2 Comparison to Cost of Introducing Collective DC Plan

In this section we examine the contribution rate that the state would have to make to a new collective DC plan, in order for taxpayers to be indifferent between the introduction of a new collective DC plan and the introduction of variable annuities. The new collective DC plan would only affect future accruals. In other words, the accumulated benefits that employees have earned based on service and salary as of today, along with associated COLAs, would be maintained in full. This is analogous to a hard freeze of the DB plan, as no new DB benefits are accrued. Furthermore, the administrative costs of managing the collective DC plan and the traditional DB plan are likely to be very similar, since the collective DC plan does not involve the creation of individual accounts or any new layers of fees.

Figure 1 shows the initial level of benefits, as a percentage of final salary, that a newly hired 25 year old worker saving 14.5% of their salary (employee plus employer contributions) might expect to receive if they retired and annuitized at age 60. The figure assumes the collective DC plan has 2/3 of its holdings in risky assets (similar to the proportion held now by DC plans). It was produced by simulating 10 million possible paths for the market, assuming a log-normal model for stocks, equity risk premia from 4-6%, and market volatility of 16%. Annuity factors were calculated, like those used in the Wisconsin PLAA, using mortality tables pooled across genders. These marginally overstate the true cost of annuities for men, and thus understate the true expected level of initial annuity payments for men, but have the opposite effect for women.

Figure 1 implies that across most paramterizations, replacement rates in the collective DC plan are not that likely to reach the levels of those guaranteed in the traditional DB plan. For example, if the

worker has a 35 year career, a traditional DB plan with a 2% benefit factor would give her a 70% replacement rate for the initial benefit. However, this reflects the fact that asset returns would have to be extremely strong going forward to meet these DB replacement rates without the kind of taxpayer contribution increases calculated in the previous section. Figure 1 also illustrates that in a collective DC plan with annuitization at market rates, the annuity a retiree can get depends very strongly on the retirement age. The lower the retirement age is, the fewer years of saving and the more years of drawing the annuity payment. Higher savings rates also allow for higher replacement rates.

We now compare the relative cost effects of introducing PLAAs versus introducing a collective DC plan for future benefits. In the PLAA case we consider that is most similar to Wisconsin, there is a 5% investment threshold and a nominal floor of the initial formula benefit. Table 7 shows that taxpayers would be indifferent from a cost perspective between introducing this kind of variable annuity and implementing a hard freeze of benefits with an employer contribution rate of 9% of pay. For the five states with the least saving potential of PLAAs relative to freezing, implementing this kind of variable annuity would only save as much as a DC plan with 19% of pay being contributed by the employer. In these instances, a collective DC with, say, a 10% contribution, would be substantially better. For the five states with the most cost saving potential of PLAAs relative to freezing, implementing the PLAAs would reduce the taxpayer burden by as much implementing a hard freeze with zero (and even negative) employer contributions.

There are several major drivers of state variation in the potential for the collective DC to reduce revenue demands on taxpayers. First, 30% of public pension system members are not in Social Security, and we assume hard freezes would entail bringing members into Social Security primarily at taxpayer expense. If a system has many members who are not in Social Security, the freeze of DB accruals in favor of the collective DC arrangement does not reduce the taxpayer burden by nearly as much as a system whose members are in Social Security. Second, if DB service accruals are particularly large, then the elimination of new DB accruals has a particularly strong downward effect on revenue demands. Third, if

employees are already paying a substantial amount into the DB plan, the freezing of DB accruals in favor of collective DC does not affect taxpayer revenue demands as strongly as if they do not.

The Social Security effect is dominant in Table 7. States such as Ohio, Maine, Illinois and Massachusetts, where most public employees are not in Social Security, generally see the PLAA implementation reducing the taxpayer burden by much more than the introduction of collective DC. In states such as New York and Pennsylvania, where employees are in Social Security, the PLAA implementation would provide a greater cost saving than the hard freeze.

6. Conclusions

This paper has considered several different paths that public sector DB systems could take towards introducing risk-sharing with employees. We have derived the costs for a menu of three possible reforms to state pension systems: 1.) the implementation of variable annuities in which benefit adjustments are linked to asset performance for all participants; 2.) permanent employee contribution increases; and 3.) the freezing of DB benefits to introduce a collective DC plan. We have found that implementing variable annuities with a nominal floor at a 5% investment hurdle would have an economic budgetary impact equivalent to requiring permanent employee contribution increases of 11.6%. Similar savings could be achieved by freezing DB accruals and introducing a collective DC plan with a 9% contribution rate.

One interpretation of these results is that employees as a group could be given an option to choose among these three reform plans, either accepting a contribution increase of 11.6% or taking one of the hybrid arrangements. One critical caveat is that we have not modeled the individual selection that would take place if each employee could choose a different option. Obviously if each individual could choose then workers closer to retirement (or already retired) would choose the hard freeze or employee contribution increase. Employees as a group would all have to choose one of the three options.

Introducing risk-sharing either though variable annuities or through a collective DC arrangement would have substantial effects on the solvency of state pension systems. The results show that achieving

meaningful reductions in the present value of accrued legacy liabilities through variable annuities would require implementation for all public employees, including those already retired, although retirees in our calculations are guaranteed that their annuities would not drop from their current levels. In a system analogous in design to that of Wisconsin, unfunded legacy liabilities nationally would be $575 billion smaller. If the variable annuities only apply to currently active workers, not retirees, the effects on total liabilities are somewhat more muted. Introducing collective DC plans instead of PLAAs would not affect the accrued legacy liability at all.

The effects on the amount of increased resources that would have to be devoted to pension systems are less dependent than the present value of the legacy liability on whether the variable annuities would apply to active workers only or to all participants including retirees. The $1,375 per household per year in additional contributions that would need to go to state systems would still fall to $872 per household per year even if it only applies prospectively for today's active participants, as opposed to $776 per household if applied to all participants including retirees. Similarly, the requisite employee contribution increases would still be almost 10% of payroll smaller (falling from 26.5% to 15.0%) even if the variable annuity only applied to today's active participants as opposed to all participants including retirees. Achieving the same revenue-demand savings as applying the PLAA to all participants including retirees could also be achieved through the collective DC plan, even though the collective DC plan does not reduce the present value of the legacy liability.

Finally, we note that while the items on the menu of reform measures we derive would each individually reduce the annual revenue demands of funding pensions by around 45% (= 1 – (776 / 1375)), they still leave a required annual revenue increase of $776 per household. This reflects the fact that a large portion of the liabilities are owed to retirees and none of the plans we consider reduces payments for current retirees.

References

Barr, N., and P. Diamond, 2008. *Reforming Pensions: Principles and Policy Choices*. Oxford University Press.

Bovenberg, L., 2011. Pension Reform in the Netherlands from an International Perspective. http://www.netspar.nl/files/Evenementen/2011-01-27%20IPW/pres%20bovenberg.pdf

Bovenberg, L., and Nijman, T., 2011. Collective Pensions and the Global Financial Crisis: The Case of the Netherlands. Working Paper.

Brown, J., Mitchell, O., Poterba, J., 2001. The role of real annuities and indexed bonds in an individual accounts retirement program. In Campbell, J.Y., and Feldstein, M. (eds.), Risk Aspects of Investment-Based Social Security Reform. Chicago, IL: University of Chicago Press, 321-360.

Brown, J., and Poterba, J., 2006. Household Demand for Variable Annuities. Tax Policy and The Economy 20, 163-192.

Brown, J., Liang, N., and Weisbenner, S., 2007. Individual account investment options and portfolio choice: Behavioral lessons from 401(k) plans 91(10), 1992-2013.

Brown, J., and Wilcox, D., 2009. Discounting State and Local Pension Liabilities. American Economic Review 99(2), 538-842.

Bulow, J., 1982. What are corporate pension liabilities? Quarterly Journal of Economics 97, 435-452.

Campbell and Viceira, 2002.

Choi, J., Laibson, D., and Madrian, B., 2011. $100 Bills on the Sidewalk: Suboptimal Investment in 401(k) Plans. Review of Economics and Statistics, forthcoming.

Clark, R., and Schieber S., 2004. Adopting Cash Balance Plans: Implications and Issues. Journal of Pension Economics and Finance 3(3), 271-295.

Coronado, J., and Copeland, P. 2012. Cash Balance Plan Conversions and the New Economy. Journal of Pension Economics and Finance 3(3), 297-314.

Gustafson, Matthew, 2012. Defined Contribution Pensions and Retirement Patterns: A Natural Experiment. Working paper.

Heaton, John, and Deborah Lucas, 1997. Market Frictions, Savings Behavior, and Portfolio Choice. Macroeconomic Dynamics, 1997 1(1), 76-101.

Horneff, W., Maurer, R., Mitchell, O., and Stamos, M., 2010. Variable Payout Annuities and Dynamic Portfolio Choice in Retirement. Journal of Pension Economics and Finance 9(2), 163-183.

Jagannathan, R., and N. R. Kocherlakota, 1996. Why Should Older People Invest Less in Stocks Than Younger People? Federal Reserve Bank of Minneapolis Quarterly Review 20, 11–23.

Novy-Marx, R., and Rauh, J., 2011a. Public Pension Promises: How Big Are They and What Are They Worth? Journal of Finance 66(4), 1207-1245.

Novy-Marx, R., and Rauh, J., 2011b. Policy Options for State Pension Systems and Their Impact on Plan Liabilities. Journal of Pension Economics and Finance 10(2), 173-194.

Novy-Marx, R., and Rauh, J., 2011c. Revenue Demands of Public Employee Pension Promises. Working Paper.

Monahan, A., 2010. Public pension plan reform: The legal framework. Education Finance and Policy 5(4), 617-646.

Milevsky, M., 2008. Portfolio Choice with Puts: Evidence from Variable Annuities. Financial Analysts Journal 64(3), 80-96.

Ponds, E., and B. van Riel, 2009. Sharing Risk: The Netherlands' New Approach to Pensions. Journal of Pension Economics and Finance 8(1), 91-105.

State of Wisconsin Investment Board, 2010. How do variable fund investments affect me? SWIB Information Paper.

Tang, N., Mitchell, O., Mottola, G., and Utkus, P., 2010. The efficiency of sponsor and participant portfolio choices in 401(k) plans. Journal of Public Economics 94, 1073-1085.

Viceira, L., 2001. Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income. Journal of Finance 56(2), 433-470.

Wisconsin Department of Employee Trust Funds, 2010. Calculating your retirement benefits. Publication ET-4107 (Rev 3/2010).

Wisconsin Legislative Audit Bureau, 2010. An Evaluation: State of Wisconsin Investment Board. Report 10-14, November 2010.

Figure 1: Distribution of possible initial annual retirement payments, as a percent of final year's salary

This figure shows the relative likelihoods of receiving different levels of annual annuity payments, for a worker that begins saving 14.5% of their salary at the age of 25, invests in a two-thirds/one-third mix of stocks and bonds, and retires at 60 purchasing an annuity with a COLA equal to half of the 2% inflation rate. Bonds are assumed to generate real returns of 1.7% (1.5% in excess of the short rate), while the equity risk premium and market volatility are assumed to be 5% and 16%, respectively. The base case, in green, is fixed across the panels, which show the impact of varying retirement age, the savings rate, the equity risk premium, and the extent of the inflation protection, respectively.

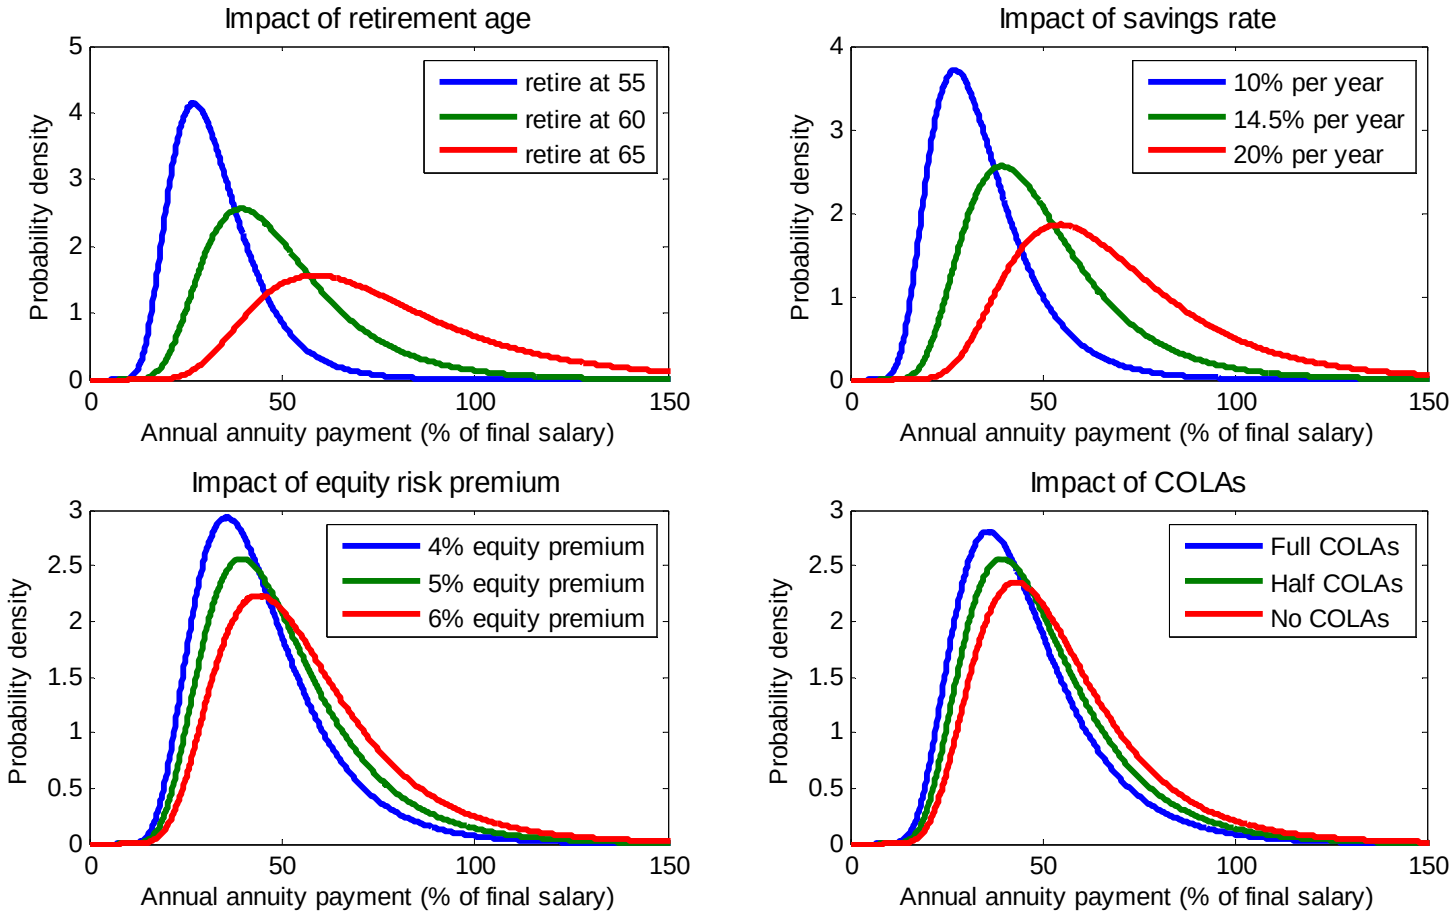


Table 1
Risk-Sharing in Various Pension Systems

	Investment Risk During Accumulation	Investment Risk During Retirement	Longevity Risk	Investment Choice
US Traditional Public Sector DB	Sponsor	Sponsor	Sponsor	Sponsor (pooled)
UK Traditional DB	Sponsor	Sponsor	Sponsor	Sponsor (pooled)
US 401(k), 403(b), Federal TSP	Participant	Participant, unless buys annuity (uncommon)	Participant chooses lump sum or (rarely) annuity	Individual
UK DC Schemes	Participant	Participant	Participant chooses lump sum or annuity	Individual (and Annuity Provider)
Cash Balance Plans: US Corporate, Nebraska PERS, Texas MRS and Texas CDRS [like collective DC with minimum return guarantee]	Participant, w/ minimum return guaranteed by sponsor	Annuity provider, and/or participant if he takes a lump sum	Participant chooses lump sum or annuity	Sponsor (pooled)
Wisconsin Retirement System (Core) [PLAA, with money purchase option in accumulation phase]	WRS guarantees formula benefit, participant can receive higher benefit if assets perform	Participant, w/ floor guaranteed by sponsor and co-insured by other participants	Participant, w/ floor guaranteed by sponsor and co-insured by other participants	Sponsor (pooled)

Alternate Plan of Certain Texas Counties (Galveston, Matagorda, Brazoria)	Participant, w/ Floor Provided by Financial Institution	Annuity provider, and/or participant if he takes a lump sum	Annuity provider, and/or participant if a he takes a lump sum	Sponsor (pooled)
Most Dutch Occupational Schemes	Shared via Conditional Indexation	Shared via Conditional Indexation	Shared	Sponsor (pooled)

Table 2

Wisconsin Retirement System Benefit Formula Parameters

	Formula Multipliers		Maximum Benefit		Normal Retirement #1		Normal Retirement #2		
Group	Pre-2000	Post-1999	Pre-2000	Post-1999	Age	With Years	Age	With Years	Early Retirement Reduction Factor
General and Teachers	1.765%	1.60%	65%	70%	65				Nonlinear Schedule
Protective Covered by SS	2.165%	2.00%		65%	53	25	54	5	Linear (4.8% p.a.)
Elected Officials	2.165%	2.00%		70%	62				Nonlinear
State Executive Plan Employees	2.165%	2.00%		70%	62				Nonlinear
Protective Not Covered by SS	2.665%	2.50%	85%	85%	53	25	54	5	Linear (4.8% p.a.)

Note: Police, prison guards, sheriffs and deputies are covered by Social Security. Firefighters are not covered by Social Security.
Source: Wisconsin Department of Employee Trust Funds (2010)

Table 3
Effect of Wisconsin Retirement System's Hybrid Features on Taxpayer Liabilities

	Present Value of Liability ($ billions)		
Performance-Linked Annuity	No	Yes	Yes
Floor	No	No	Yes
ABO	106.0	80.7	94.2
EAN	119.3	90.9	106.1
PVB	136.4	103.9	121.3

Note: The table shows the total present value liability in trillions of dollars as of December 2010 for the Wisconsin Retirement System (WRS), with deflated cash flows discounted at the TIPS yield curve. Liabilities are stated on three standards: Accumulated Benefit Obligation (ABO), Entry Age Normal (EAN), and Present Value of Benefits (PVB). ABO is the narrowest measure, reflecting only benefits based on service and salary earned up to today. The left-hand column shows the liability without performance-linked annuity adjustments (PLAAs) or floors. The middle column shows the liability accounting for the performance-linked annuity adjustments (PLAAs) without a floor. The right column shows analogous figures for the Wisconsin PLAAs with a nominal floor of the initial formula benefit, as the WRS provides.

Table 4

Liability Effects of Implementing Variable Annuities in All 50 States, $ Trillions

		PLAAs Without Floor Applied to					
	Status Quo	Actives Only			All Participants		
Hurdle	NA	3%	5%	7%	3%	5%	7%
ABO	4.63	4.38	4.07	3.85	4.14	3.44	2.94
EAN	5.10	4.80	4.42	4.14	4.56	3.79	3.22
PVB	6.77	6.25	5.60	5.13	6.01	4.97	4.22

		PLAAs With Floor Applied to					
	Status Quo	Actives Only			All Participants		
Hurdle	NA	3%	5%	7%	3%	5%	7%
ABO	4.63	4.53	4.33	4.24	4.50	4.05	3.85
EAN	5.10	4.98	4.73	4.62	4.95	4.46	4.23
PVB	6.77	6.56	6.14	5.95	6.53	5.86	5.56

Note: The table shows the total present value liability in trillions of dollars as of December 2010, with deflated cash flows discounted at the TIPS yield curve. Hurdle rates are threshold rates for annuity adjustment (not discount rates). Liabilities are stated on three standards: Accumulated Benefit Obligation (ABO), Entry Age Normal (EAN), and Present Value of Benefits (PVB). ABO is the narrowest measure, reflecting only benefits based on service and salary earned up to today. The top panel shows the liability if performance-linked annuity adjustments (PLAAs) were applied without floors to actives only (middle panel) and all panels (right panel) respectively. The lower panel shows analogous figures if the PLAAs are implemented with a nominal floor of the initial formula benefit.

Table 5

Full Funding Revenue Demands per Household and Effects of Introducing Variable Annuities

	All States	Highest 5 States	Lowest 5 States
Status Quo	$ 1,375	$ 2,154	$ 549
PLAAs without Floor			
Actives Only			
3%	1,155	1,669	472
5%	798	1195	256
7%	543	856	101
All Participants			
3%	1,069	1,474	444
5%	605	857	158
7%	270	413	-48
PLAAs with Floor			
Actives Only			
3%	1,198	1,719	503
5%	872	1,282	309
7%	655	989	181
All Participants			
3%	1,170	1,602	512
5%	776	1,075	274
7%	528	740	126

Note: The Status Quo line shows the amount of increased revenue per US household per year that would be required to achieve full funding of state pension systems in 30 years on an ABO basis. Methodology is as in Novy-Marx and Rauh (2011c), although here we present figures for state plans only. The top panel shows the effects of introducing PLAAs without any floor under benefits. The bottom panel shows the effects of introducing PLAAs with a nominal floor of the initial formula benefit. The left column shows the averages for all states. The middle column shows the averages for the five states with the highest per-household revenue demands (New Jersey, Oregon, Rhode Island, Ohio, and Wyoming). The right column shows the averages for the states with the lowest per-household revenue demands (West Virginia, Arkansas, Arizona, Utah, and Indiana). The scenario in the box is the one closest to the Wisconsin Retirement System's design.

Table 6

Full Funding Revenue Demands as Percent of Payroll and Effects of Introducing Variable Annuities

	All States		Highest 5 States		Lowest 5 States	
	Required Increase	Savings vs. Status Quo	Required Increase	Savings vs. Status Quo	Required Increase	Savings vs. Status Quo
Status Quo	26.5%		39.9%		12.3%	
PLAAs without Floor						
Actives Only						
3%	22.3%	4.2%	30.9%	8.9%	10.1%	2.2%
5%	15.4%	11.1%	23.7%	16.2%	5.1%	7.2%
7%	10.5%	16.0%	18.5%	21.3%	1.6%	10.7%
All Participants						
3%	20.6%	5.9%	25.6%	14.2%	9.5%	2.7%
5%	11.7%	14.9%	15.5%	24.3%	3.1%	9.2%
7%	5.2%	21.3%	8.3%	31.6%	-1.5%	13.8%
PLAAs with Floor						
Actives Only						
3%	23.1%	3.4%	32.0%	7.8%	10.7%	1.5%
5%	16.8%	9.7%	25.6%	14.3%	6.3%	6.0%
7%	12.6%	13.9%	21.4%	18.4%	3.4%	8.9%
All Participants						
3%	22.6%	4.0%	28.3%	11.6%	11.0%	1.3%
5%	15.0%	11.6%	20.0%	19.8%	5.6%	6.6%
7%	10.2%	16.3%	15.0%	24.8%	2.3%	10.0%

Note: Status Quo is the increased revenue as a share of public employee payroll required to achieve full funding of state pension systems in 30 years on an ABO basis. Methodology is as in Novy-Marx and Rauh (2011c), although here we present figures for state plans only. The left panel shows the averages for all states. The middle panel shows the averages for the five states with the highest revenue demands as a share of payroll (Colorado, Pennsylvania, Michigan, Illinois, Oregon). The right panel shows the averages for the states with the lowest revenue demands as a share of payroll (Idaho, Hawaii, Utah, Arizona, Indiana). The scenario in the box is the one closest to the Wisconsin Retirement System's design. In each panel there are two columns: the required contribution as a share of payroll to achieve full funding in 30 years (Required Increase), and the savings achieved for the PLAA implementation relative to the status quo.

Table 7

Employer Contribution Rate for Equivalent Cost Saving from Collective DC Introduction

	All States	5 States with Least Saving Relative to Freeze	5 States with Greatest Saving Relative to Freeze
PLAAs without Floor			
Actives Only			
3%	16%	25%	7%
5%	10%	17%	0%
7%	5%	12%	-6%
All Participants			
3%	15%	25%	4%
5%	6%	16%	-7%
7%	-1%	10%	-14%
PLAAs with Floor			
Actives Only			
3%	17%	26%	8%
5%	11%	19%	2%
7%	7%	14%	-3%
All Participants			
3%	17%	27%	7%
5%	9%	19%	-2%
7%	4%	14%	-7%

Note: The table shows the employer contribution rate for new defined contribution plans that would make state taxpayers indifferent between implementing variable annuities and introducing a new collective DC plan with a hard freeze of all DB accruals. The scenario in the box is the one whose costs would most closely match those of switching to the Wisconsin Retirement System's variable annuity design.