UNIVERSITY OF ARKANSAS | College of Education & Health Professions
Education Reform

WORKING PAPER SERIES

Economics of Sustainable Public Pension Funding

Robert M. Costrell and Josh McGee

6-10-2022

EDRE Working Paper 2022-07

ECONOMICS OF SUSTAINABLE PUBLIC PENSION FUNDING

March 13, 2022, v3.3

Robert M. Costrell and Josh McGee
Department of Education Reform
200 Graduate Education Building
University of Arkansas, Fayetteville, AR 72701
costrell@uark.edu joshmcgee@uark.edu

Prepared for 47th Annual Conference of Association for Education Finance and Policy (AEFP),
Denver CO, March 17, 2022

ABSTRACT: In this paper we propose a new approach to sustainable public pension funding, as an alternative to: (i) traditional actuarial full-funding policies, on the one hand; and (ii) recent proposals aimed instead at stabilizing pension debt at current levels. Actuarial contribution policies aim to fund liabilities that are wrongly discounted at the expected rate of return on risky assets; and these policies promise to do so with amortization schedules that terminate in a precipitous future drop in contributions, which never materializes. Conversely, recent debt-stabilization proposals (Lenney, Lutz, and Sheiner, 2019a; 2019b) properly discount liabilities at a risk-free rate, but effectively untether contribution policy from those liabilities. Our analysis integrates properly discounted liabilities with investment strategies that may be risk-tolerant to some degree, in a policy framework that more transparently conveys the tradeoffs we face.

We begin with the fundamental equations of motion for assets and liabilities – how these two sides of the ledger evolve with contributions, asset returns, and newly accrued liabilities. From these equations we formally derive the characteristics of steady-state pension funding – which we take as the definition of sustainability. We also derive the set of contribution adjustment parameters that smoothly achieve steady-state – a non-trivial exercise. The resulting contribution schedules differ conceptually from the traditional setup of normal cost plus amortization. Building on previous work (Costrell, 2018, Costrell and McGee, 2020), we examine the steady-state implications of differentiating between the assumed return on assets (r) and the discount rate on liabilities (d). We integrate these insights into a semi-formal social optimization framework to sketch out a contribution policy approach that conveys the tradeoffs between intergenerational burden-sharing, the pursuit of returns, and the cost of risk-bearing.

KEYWORDS: pension finance

JEL CODE: I22, H75

ACKNOWLEDGMENTS:

ECONOMICS OF SUSTAINABLE PUBLIC PENSION FUNDING

Robert M. Costrell and Josh McGee
University of Arkansas
March 13, 2022, v3.3

Introduction and Summary

Sustainability is the central concern of public pension funding. But the precise meaning of the term, let alone the conditions for its satisfaction, vary with the user and is often not precisely defined or analyzed at all. In this paper, we propose a formal economic definition of sustainability as a funding policy that generates a steady state in the contribution rate and funded ratio. We begin with the standard equations of motion for assets and liabilities, contingent on the parameters of the system and the funding policy, which govern the trajectory of contributions. From these parameters (rate of return on assets, discount rate for liabilities, payroll growth rate, normal cost rate, benefit payout rate) and the policy specification, we derive the steady state contribution rate. This allows us to analyze the determinants of contributions in a formal model and compare that rate with actual contributions currently observed in public pension plans.

Although this exercise oversimplifies actual systems, since the parameters themselves never settle into steady states, the approach offers insights akin to other simple economic models. Steady state analysis lays out the system's resting point, even if it is a moving one. That said, there are two further important issues to explore. The first is the issue of convergence. Steady states are of less interest if the system does not converge. We find that the conditions for convergence/non-convergence in such systems are surprisingly non-trivial, even in this simple class of models. We analyze this issue formally, determining the range of adjustment parameters that yields convergence. This generates a novel contribution policy outside of steady state, differing markedly from the standard actuarial policy of normal cost plus amortization.

The second issue is how to determine the target funded ratio in the presence of risk. Specifically, our key steady-state result shows how -- depending on the target funded ratio -- the contribution rate depends on the gap between the assumed rate of return on risky assets and the low-risk discount rate on guaranteed benefits for properly valuing liabilities. Specifically, the steady-state contribution rate can fall well below the normal cost rate, due to implicitly assumed arbitrage profits represented by that gap. This leads us to sketch out a contribution policy approach that conveys the tradeoffs between intergenerational burden-sharing, the pursuit of returns and risk tolerance.

Operationally Defining Sustainability

Although pension plan sustainability is a central concept in policy discussions, the usage of the term varies and it is not always well-defined. In general terms, the underlying question is whether the current plan can continue more or less as is, or whether it will require substantial change (such as a rise in contributions) to stave off insolvency or some other form of collapse. We believe the formalization of this concept is to be found in steady-state analysis. What would a steady-state look like under current plan parameters, provisions, and policies, and how does the contribution rate in such a steady state compare with current rates? If the steady-state contribution rate is significantly higher than the current rate, then one might well conclude the system is not sustainable with current contributions. This framework still leaves several issues to be specified as we continue with this analysis, but it seems a good starting point at least. Steady-state analysis is based on the fundamental laws of motion of a pension plan, those of assets and liabilities. Let us begin with assets.

Steady State Condition for Contributions and Asset Accumulation

We can consider asset accumulation as either a stand-alone basis for ascertaining sustainable funding policies, or as the first step toward considering policies that set asset targets tied to liabilities. In this section we consider asset accumulation on a stand-alone basis, which generates insights of its own, and then bring in liabilities in the next section.

There are two sources of pension funding and two uses: contributions and investment income go to cover the payment of benefits and the accumulation of assets. Of these four flow variables, the stream of benefit payments is exogenous to our analysis (determined by the tiered benefit formulas and workforce assumptions), and investment income is governed by the sequentially determined stock of assets and the exogenous series of annual returns.[1] This leaves the series of contributions and that of asset accumulation, which are mechanically linked. That is, the *funding policy is simultaneously a contribution policy and an asset accumulation policy*.

Formally, this relationship is captured in the fundamental asset growth equation:

$$(1)\ A_{t+1} = A_t(1+r_t) + \boldsymbol{c_t} W_t - c^p{}_t W_t\,,$$

where A_t denotes assets at the beginning of period t, r_t is the return in period t, W_t is payroll, while $\boldsymbol{c_t}$ and $c^p{}_t$ are the contribution and benefit payment rates, respectively, as proportions of payroll (Table 1 lists notation). Assets grow by investment earnings, plus contributions, net of benefit payments. Equation (1) is simply an accounting identity. To give it economic content, for sustainability analysis, we need to specify a funding policy to drive $\boldsymbol{c_t}$. Given returns and benefit payments, the contribution policy sets asset growth. We will spell out our approach to the choice of contribution policy below, but even before doing so, equation (1) helps focus on the fundamental tradeoffs among these policies without getting overly distracted by their details.

[1] Of course, the exogeneity assumed here is conditional on the investment policy, i.e., the asset allocation.

It will be useful to re-express equation (1) in terms of the ratio of assets to payroll, $a \equiv (A/W)$. Dividing through (1) by W_t, and denoting the growth rate of payroll by g, we have:

$$(1')\ a_{t+1}(1+g_t) = a_t(1+r_t) + \mathbf{c_t} - c^p_t\,.$$

The big picture here can be illuminated by examining the steady-state relationship between contributions and assets. In steady-state, the growth of assets must equal the growth of payroll, so the asset ratio is constant, $a_{t+1} = a_t = a^*$. Removing the time subscript for the steady-state values of the benefit payment rate c^p, the rate of return r, and the payroll growth rate g, we have the relationship between the steady-state values of the contribution rate and the asset ratio:

$$(1^*)\ \mathbf{c^*} = c^p - (r - g)a^*.$$

The interpretation is straight-forward: benefit payments are covered by a mix of contributions and investment income (net of growth), where the mix is determined by the funding policy. Under a policy of pay-go, where no assets are accumulated ($a^* = 0$), the contribution rate must cover the benefits payment rate c^p. Under a policy of pre-funding, to one degree or another, the goal is to accumulate a certain asset level, a^*, so the income from those assets (net of growth) can help fund benefits, ultimately reducing reliance on contributions.

One very simple test of sustainability is to consider whether current contribution rates are sufficient to sustain a steady state at current asset levels. That is, if we set $a^* = a_0$, would the current contribution rate, c_0, need to rise or not to sustain the asset level?

Let us consider the trends and magnitudes of the relevant variables. Figure 1 depicts the aggregate values of $\mathbf{c_t}$ and c^p_t for FY01 – FY20, of the 119 state and 91 local plans in the Boston College Public Plans Data, which account for 95 percent of state and local pension assets and members in the U.S. As is well-known, the contribution rate, as a percent of payroll, has been steadily climbing since the turn of the century, from about 12 percent to 27 percent. The benefit (or "pay-go") rate has also trended up, from 20 percent, but may now be leveling off at about 38 percent.[2] It is important to note that throughout this period the benefit rate exceeds the contribution rate by a large margin, exceeding 10 percentage points since 2010. That is, the basic cash flow (excluding investment income) is negative, due to some combination of plan maturity and possibly some contribution shortfall (the question we are considering in some form). Thus, if assets were to be depleted, contributions would have to jump to cover benefits.

Figure 2 depicts the asset ratio $a \equiv (A/W)$ from the same dataset. This has fluctuated with market returns and has also been affected by the trends in benefit payments, but in recent years assets have hovered around a multiple of 5 times covered payroll.

For illustrative purposes, we can consider typical plan assumptions of $g = 3\%$ and $r = 7\%$ to calculate $c^* = c^p - (r - g)a_0 = 0.38 - (0.07 - 0.03) \times 5 = 0.18 < c_0 = 0.27$. Thus, taken at face value, this would suggest that, in the aggregate, the current configuration is not only sustainable, but that contribution rates could fall and still support current asset ratios. Of course, this depends on a host of assumptions, not least of which are the assumed rate of return and growth rate. However, we can see that as long as $(r - g)$ exceeds about 2 percent (e.g. $r > 5$ percent), current contributions could be sustainable in the aggregate.

This picture also holds generally for the individual plans in the PPD database. Using each plan's assumed return (the vast majority lie between 7.0 and 7.5 percent for FY20), we find that in 158 of the 188 plans for which c^* can be calculated, the contribution rate exceeds that value.[3] This also holds for 69 of the 79 largest plans, with assets exceeding $10 billion. Reducing each plans' assumed return to 5.0 percent changes the picture. Under this calculation,

---

[2] Lenney, Lutz, and Sheiner (2019a; 2019b) project that the benefit rate will peak around [year?] and decline thereafter, as recent hires, in less generous tiers, enter retirement, and beneficiaries of more generous tiers die.

[3] The assumed growth rate for payroll is only available in the PPD for 76 plans. Of those, the vast majority lie between 2.75 and 3.5 percent, so we set the growth rate at 3.0 percent for the calculation of c^* in all plans.

the contribution rate for most plans (107 of the 188 plans, and 48 of the largest 79 plans) is not high enough to sustain the current asset ratio.

Contribution Policy for Convergence to Steady State Asset Accumulation: Analytics

Although steady-state calculations (those above, and further specified below) are instructive, they are not compelling unless there is a dynamic process that converges toward a steady state. Of course, the steady state is always a moving target, as the parameters c^p, r, and g vary over time, but we can analyze whether the system moves in the right direction at any given time, taking these parameters as constants, at their steady state values.

Convergence is not automatically assured, as can be discerned by considering the asset accumulation equation (1′) alone (before adding in a contribution policy equation). To simplify notation, let $R = 1+r$, $G = 1+g$, and re-express (1′) as:

$$(1'')\ a_{t+1} = a_t(R/G) + (\boldsymbol{c_t} - c^p)/G.$$

For $R > G$ (as usually assumed), the coefficient on the prior value of the state variable a exceeds one, which is destabilizing. For example, suppose we consider a policy that sets the contribution rate to some target rate and holds it constant.[4] Unless that target rate corresponds to the steady-state value for maintaining the *current* asset ratio, the system will diverge. Stated alternatively, suppose one aims at an asset ratio $a^* \neq a_0$, and immediately sets $c = c^*$ (using (1*)), jumping up or down from $\boldsymbol{c_0}$, and holding it there. Then the system will move away from a^*, rather than toward it. If a^* is set greater than a_0, then a_t, will shrink further away from a^*, and conversely if a^* is set lower than a_0.[5] The reason is straightforward. Setting a higher a^* means setting a lower

---

[4] This is not a fanciful policy scenario. The Lenney, Lutz, and Sheiner (2019a; 2019b) policy simulation is to set c equal to a steady-state value and hold it there.

[5] Formally, the solution is $a_t = a^* + (R/G)^t(a_0 - a^*)$.

c^* for $r > g$ (see equation (1*)), since one expects to rely on higher investment income, in lieu of contributions, to cover benefits. But since assets are not yet at that higher level of a^*, the investment income falls short of that which would obtain in the steady state one aspires to. Thus, by prematurely setting contributions at the correspondingly low level c^* one embarks on a path of asset decumulation. And conversely for $a^* > a_0$.

So, what would a contribution policy look like that converges to a steady state targeted at a^* with contributions c^*? It might be thought that an adjustment process that gradually closes the gap between current contributions and c^*, rather than a sudden jump to c^* would do the job, but as we shall see below, it will not. The reason, as would be suggested by the discussion above, is that the contribution required to cover benefits depends on the gap between current assets and a^*. Alternatively, one might then suppose that an adjustment process for contributions based on the *asset gap* would do the job. However, as we shall see, that will not suffice either. For a convergent path, we show that the policy should adjust contributions based on both gaps, between c^* and c_t and between a^* and a_t, in combinations to be derived below.

Before doing so, note that the policy we are deriving differs not only from a discrete jump to c^*, but also from the trajectory of actuarial funding policies. The actuarial payment schedule is either a constant percent of payroll, or ramps up to such a rate, and then falls off a cliff at the end of the amortization period, once full funding is expected to be achieved. The policy we derive below aims to converge on a steady state, either monotonically or through dampened oscillations (depending on the adjustment speeds chosen), and then to stay there.

Specifically, consider a contribution and asset-accumulation policy that starts by specifying a target asset ratio, a^* (more on how that might be chosen, in a later section), then calculates the corresponding steady-state contribution rate c^*, using (1*) above. We then posit a contribution policy that adjusts the contribution rate based on the gaps between the target and actual asset ratio and contribution rate:

(2) $c_{t+1} = c_t + \beta(c^* - c_t) + \gamma(a^* - a_t)$, where $\beta \epsilon (0,1)$.

Together with (1″), we have a simple system of two linear difference equations that can be usefully expressed in matrix form:

$$\begin{bmatrix} a \\ c \end{bmatrix}_{t+1} = \begin{bmatrix} (R/G) & (1/G) \\ -\gamma & (1-\beta) \end{bmatrix} \begin{bmatrix} a \\ c \end{bmatrix}_t + \begin{bmatrix} (-c^p/G) \\ (\gamma a^* + \beta c^*) \end{bmatrix}.$$

Denote the transition matrix above by A. Then the asymptotic stability condition is[6] $|tr\,(A)| < 1 + det(A) < 2$, which, in the present case, implies

(i) $\gamma > \beta(R - G) \equiv \gamma_{min} > 0$, and

(ii) $\gamma < G - R(1 - \beta) \equiv \gamma_{max}$.

Condition (i) shows formally what was alluded to above: a piece of the adjustment mechanism must be based on the asset gap, not just that of the contribution rate. The logic is straight-forward. Suppose the contribution rate is already at its target c^*, but the asset level is below the target a^*. Then contributions will have to rise in the short run to accumulate more assets, before eventually dropping back down toward c^*. Condition (ii) implies that the adjustment mechanism must include the contribution gap, too. Formally, since we must have $\gamma_{max} > \gamma_{min}$, this requires $\beta > (R - G)/G > 0$. The logic here is also straight-forward. If assets are at their target ratio, but the contribution rate is below c^*, then it needs to rise.

As our discussion above suggests, the convergence to steady-state may not be monotonic. Indeed, it may not only reverse direction once (asymptotically monotonic), it may be oscillatory. The condition for asymptotic oscillation is $[tr\,(A)]^2 < 4 \cdot det(A)$, or, in the present case:

[6] See, for example, Neusser (2021), equation (3.18), p. 84.

(iii) $\gamma > G[(R/G) - (1-\beta)]^2/4 \equiv \gamma_{m/o},$

where the subscript *m/o* denotes the border between monotonic and oscillatory. It can be shown that for $\gamma_{max} > \gamma_{min}$ (i.e., $\beta > (R-G)/G$), $\gamma_{m/0}$ lies in between. Thus, the asymptotic behavior of the system varies with the range of γ as follows:

Table 2: Convergence Conditions	
Range of γ	**Asymptotic Behavior of *(1″)-(2)***
$\gamma < \gamma_{min}$ (given by (i))	Monotonic divergence
$\gamma_{min} < \gamma < \gamma_{m/o}$ (given by (iii))	Monotonic convergence
$\gamma_{m/o} < \gamma < \gamma_{max}$ (given by (ii))	Oscillatory convergence
$\gamma_{max} < \gamma$	Oscillatory divergence

Figure 3 illustrates the combinations of β and γ that correspond to these asymptotic behaviors. In general, it seems reasonable to presume that policy-makers would prefer monotonic convergence to oscillatory convergence. Thus, the combinations of β and γ to be considered would lie between γ_{min} and $\gamma_{m/o}$, depicted by the black and blue curves in Figure 3.

Contribution Paths Toward Steady State Asset Accumulation: Simulation

Armed with these analytics, we illustrate some dynamic paths for contributions and assets under policies that might plausibly be suggested. We begin with the representative plan assumptions given above, $R = 1.07$, $G = 1.03$, $c^p = 0.38$, $c_0 = 0.27$ and $a_0 = 5$. If, as discussed in the previous section, our goal is simply to maintain the current asset level, then $a^* = 5$ and $c^* = 0.18$. In this case, policy-makers might be expected to choose a path that reduces contributions as quickly as possible, without overshooting. This means choosing the adjustment parameters β

and γ arbitrarily close to one and γ_{min} respectively, at the right boundary of the black γ_{min} curve in Figure 3. Under these parameters ($\beta = 1.0$, $\gamma = \gamma_{min} = 0.04$), simulation confirms the contribution rate drops immediately to just below c^*, while the asset ratio shades slightly above $a_0 = a^*$.

Suppose we consider a more ambitious target ratio of $a^* = 7$. This increase of 40 percent above a_0 would accumulate approximately the assets needed to match liabilities (discussed in the next section), i.e., full actuarial funding (discounted at the expected return). At $a^* = 7$, (1*) gives us $c^* = 0.38 - (0.07 - 0.03) \times 7 = 0.10 < c_0 = 0.27$, thus allowing eventually for a dramatically lower contribution rate. Here, the choice of adjustment parameters β and γ must navigate an intertemporal policy tradeoff. Contributions need to rise in the short run to accumulate the assets required for the long-term reduction to c^*. Thus, the tradeoff is between speed of reaching c^* vs. tempering the short-term rise in c required to reach a^*. Suppose we set a target of approaching c^* by year 30 (corresponding to a somewhat conventional time horizon for actuarial amortization schedules) and set the contribution adjustment parameter β equal to 0.5 (half speed). Then we find that the tradeoffs are plausibly managed by choosing the asset adjustment parameter γ near the maximum value for monotonic convergence, $\gamma_{m/o} = 0.075$. Figure 4a depicts the corresponding paths for the contribution rate (red curve, on the right scale) and asset ratio (blue curve, on the left scale). This path raises the contribution rate for about 7 years to a maximum of 36 percent (a 9 point hike), before ultimately dropping down to approximately 10 percent by year 30. Setting β any faster requires a sharper short-term rise in contributions and setting it any slower fails to so closely approach c^* in 30 years.

Naturally, these results are sensitive to the assumed rate of return. As discussed above, if $r = 5\%$ instead of 7%, the required contributions are rather different. Under the first scenario, to simply maintain the current asset ratio of 5, the contribution rate must rise by a point, instead of falling by 9 points. Similarly, for the more ambitious scenario of raising a^* by 40% (to reach full funding), the picture is more daunting if $r = 5\%$. As Figure 4b shows, the contribution rate would need to rise in the short run by 20 percentage points (to nearly 50 percent of payroll), for an ultimate reduction of only about 3 points, a much more challenging picture than Figure 4a.

There are several take-aways from these exercises. First, our dynamic analysis shows how to generate smooth adjustment paths, unlike the actuarial scenario of the contribution cliff that is supposedly reached upon completion of the amortization schedule. Second, as is well-known, but illustrated here in a formal dynamic context, deterministic scenarios have their limitations, given the risk of investment returns. Finally, even within a deterministic context, one needs some criterion to anchor the asset accumulation goal. That criterion has traditionally been based on liabilities, to which we now turn.

Steady State Condition for Liabilities

We begin with the fundamental growth equation for liabilities:

$$(3)\ L_{t+1} = L_t(1+d) + c^n{}_t W_t - c^p{}_t W_t\,,$$

where L_t denotes accrued liabilities at the beginning of period t, d is the discount rate, and $c^n{}_t$ is the "normal cost rate," the rate at which new liabilities accrue, as a percent of payroll. Liabilities grow by the interest on past liabilities, plus newly accrued liabilities, net of benefit payments that extinguish prior liabilities. Equation (3) is analogous to the asset growth equation (1), but with some key differences:

First, the formulation in (3) allows for a distinction between the discount rate d and the rate of return on assets r. Standard actuarial practice, of course, has traditionally equated the two. By contrast (as is well known and much-discussed), finance economics has consistently made the case that guaranteed benefits should be discounted by interest rates of correspondingly low-risk bonds, at least for accounting purposes. If asset accumulation, and projections thereof, continue to reflect actual and assumed returns on a higher-risk pension fund portfolio, this raises the question of how a dual rate system should play out in contribution policy. In the previous section, where our analysis was confined to asset accumulation, the contribution policy, both in steady-state and in adjustment to steady-state, depended only on r and not on d. We consider below how the consideration of liabilities, discounted at $d < r$, should or should not factor into contribution policy.

The second difference between the liability growth equation (3) and the asset accumulation equation (1) is the role of c^n_t, the normal cost rate, vs. c_t, the contribution rate. The normal cost rate is determined independently of the contribution policy. It is completely driven by the benefit formula, the cohort's assumed separation probabilities over its members' careers, and the discount rate.[7] The normal cost rate may be used to help determine the contribution policy (as in standard actuarially determined contributions), but if the benefit formula is taken as exogenous to our analysis, equation (3) is stand-alone. It is recursively prior to the asset accumulation and contribution equations. We will return to this point below.

To examine the dynamics of liability accrual, we express (3) in the state variable $\lambda = L/W$ = liabilities/payroll, using the same steps as in the derivation of (1′):

$$(3') \ \lambda_{t+1}(1+g_t) = \lambda_t(1+d) + c^n_t - c^p_t\,.$$

If we take the benefit formula and demographic/worklife assumptions as exogenous, then so are c^n and c^p. Thus, we can readily derive the steady-state liability ratio:

$$(3^*) \ \lambda^* = (c^p - c^n)/(d - g).$$

[7] It also depends on the specific actuarial cost method for allocating liabilities between past and future accruals. To fix ideas, we have in mind the standard entry age normal cost method.

This expression has a simple interpretation. First note that the present value of future payroll in steady-state is $W_t/(d - g)$, consistent with the standard formula for a growing perpetuity. Then note that the present value of future benefit payments and future liability accruals (normal costs) are, respectively, fractions c^p and c^n of the PV of future payroll. Thus, equation (3*)'s steady-state ratio between accrued liabilities and payroll represents the difference between the present values of future benefit payments and future normal costs (scaled to current payroll).[8] A decrease in d raises the former more than the latter, since future benefit payments for any given cohort (and thus for all cohorts taken together) have longer duration than future normal cost payments. Thus, λ^* rises.

It is worth clarifying here that (3*) *must* hold, as an accounting identity, if we are in demographic steady-state, with a constant growth rate g (along with unchanging separation probabilities and benefit formula). Any deviations of the liability ratio from the steady-state value can only be due to past or future variation in payroll growth, in the plan's run-up to (or run-down from) the mature membership configuration of steady-state age distribution among actives and retireds. In such non-steady-state periods, the payment rate, $c^p{}_t$, would deviate from the steady-state value c^p (lower in the run-up to plan maturity, due to lower ratio of retireds/actives, and conversely in the run-down from maturity) and that would drive the deviations of λ_t from λ^* through the accounting identities of (3) and (3′). Thus, unlike the asset accumulation dynamic, where deviations from steady-state arise from the contribution history, and which pose a non-trivial question of stability, as examined above, there is no such issue here:

[8] This follows from the basic identity that the present value of all future benefit payments equals the present value of benefits yet to be accrued (the present value of future normal costs) plus the present value of benefits previously accrued, but not yet paid out. The latter term is the accrued liability, so it equals the difference between the present value of all future benefits and the present value of future normal costs.

non-steady-state liability ratios resolve themselves simply by virtue of evolving benefit payment rates in the transition to the steady-state demographic configuration.

Figure 5 depicts the aggregate liability ratio, drawing again on the Public Plan Database, where the liabilities are reported based on each plan's assumed return, r. That ratio (depicted by the red curve) has gradually risen from about 4.6 in FY01 to about 7.2 in FY20. Several factors have contributed to this trend, including reductions in the assumed return and a rise in the ratio of retireds to actives, as plans have matured and gone beyond maturity.[9,10] Liabilities are much higher when discounted at a low-risk rate d, instead of r. Estimates vary, comparing the liability estimates of the Federal Reserve Board of Governors (depicted by the blue curve in Figure 5)[11] with those of the PPD suggest that properly discounted liabilities are 60 percent higher.[12]

Linking Asset Accumulation and Contributions to Liabilities

The natural link between our steady-state analysis of asset accumulation and liabilities is to tie the asset goal to liabilities. Of course, this is the actuarial goal of full funding. We here consider the more general goal of a target funded ratio, f^* (e.g., the putative "standard" of 80 percent funded).[13] Setting the asset goal of $a^* = f^*\lambda^*$, and, for the moment, following the actuarial convention of $d = r$, we find, from (1*) and (3*):

$$(4)\ \mathbf{c^*} = c^p - (r - g)f^*\lambda^* = c^p - f^*(c^p - c^n) = (1 - f^*)c^p + f^*c^n.$$

[9] Benefit changes of course, have also affected the trends, but in no simple fashion, as many plans raised benefits in the early 2000's and then cut benefits for new hires in the 2010's.
[10] Comparing the liability ratios with the calculated values of $(c^p - c^n)/(r - g)$ for FY01, FY10, and FY20, we find these values match for FY01 (4.6 vs. 4.5), but for FY10 and FY20, the liability ratios exceed the calculated values, 5.7 vs. 4.3 and 7.2 vs. 6.1, respectively. There are many potential explanations for these gaps, but they would be consistent with plans that are beyond mature, rather than in steady state.
[11] The denominator in the ratio depicted is the PPD payroll series.
[12] Estimates from Lenny, Lutz, Schule, and Sheiner (2020), using reported liabilities and rediscounted liabilities indicate that the latter is about 80 percent higher.
[13] See Costrell, 2018, where equation (4) was previously derived.

As the funded goal varies from zero to full funding, the steady-state contribution rate varies from the pay-go rate to the normal cost rate, with a weighted average of the two for intermediate funding targets.

Let us now consider the steady-state implications of a dual rate system: discount rate d for liabilities and assumed return r on assets. We then have:

$$(4')\ \mathbf{c^*} = c^p - (r-g)f^*\lambda^* = c^p - [(r-g)/(d-g)]f^*(c^p - c^n).$$

As before, if the funding goal f^* is zero, the contribution target is pay-go, and as f^* is set higher, the contribution target falls.

However, our question here is the impact on c^* of reducing d below r. We have already seen from (1*) that the only avenue for a drop in d to affect c^* is through its impact on the asset target a^*. Since we are considering asset goals of the form $a^* = f^*\lambda^*$, this means that a drop in d below r would raise the target contribution rate through a rise in the liability ratio λ^* unless it is offset by a reduction in the target funded ratio f^*.

If, for example, we take as our funding goal to simply maintain the current asset ratio, $a^* = a_0$, then the rise in λ^* from revaluation at d would, in effect, be completely offset by an implicit drop in the target funded ratio f^*.[14] In this polar case, setting d to a low-risk rate for the valuation of liabilities is purely an accounting and reporting measure, unrelated to funding goals.

More generally, however, one might expect that recognizing liabilities as guaranteed and that asset returns are not, might lead policymakers to consider how much risk they wish to bear and how much they wish to defray with higher contributions. To help elucidate the issue, let us consider further the steady-state contribution rate c^*, given in (4'). Here, we must interpret r as the expected return on assets (rather than a deterministic rate) that exceeds the risk-free rate d

[14] This is implicit in the Lenney, Lutz, and Sheiner (2019a; 2019b) model. This explains why the contribution rate in their model is effectively independent of d, despite their claim that setting $d < r$ makes their model conservative.

used for evaluating liabilities and their rate of accrual c^n. The first implication of this is that a full-funding policy $f^* = 1$, or anywhere near it, implies that $c^* < c^n$: the contribution rate will not cover normal costs (properly evaluated). Formally, *(4')* implies

$$(4'')\ \mathbf{c^*} - c^n = (c^p - c^n)(1 - [(r - g)/(d - g)]f^*) < 0, \text{ for } f^* > [(d - g)/(r - g)].$$

To fix magnitudes here, consider the values we have been using, $r = 0.07$ and $g = 0.03$, along with $d = 0.04$ (a typical discount rate used in private pension accounting). The critical value of f^* in the expression above is then 25 percent. For any target funded ratio exceeding 25 percent, steady-state contributions need not cover the normal costs (when rediscounted at d). Note how strikingly this contrasts with standard actuarial funding schedules, under which contribution rates drop to (but not below) c^n, upon reaching full funding.

The point can be illuminated by re-writing *(4')* and simplifying to obtain:

$$(4''')\ \mathbf{c^*} = c^p - (d - g)f^*\lambda^* - (r - d)f^*\lambda^* = (1 - f^*)c^p + f^*c^n - (r - d)f^*\lambda^*.$$

Comparing with (4), we have a rediscounted normal cost rate (higher c^n), but the third term may be interpreted as the implicitly assumed arbitrage profits between the return on accumulated assets and interest on covered liabilities. These assumed arbitrage profits help defray the higher normal costs, in lieu of contributions that might otherwise be required. Alternatively, this term may be interpreted as the risk premium, which would be borne by the plan as the implicit cost of risk under the contribution policy implied by this approach.

Sketching Out an Approach to Integrating Dual Rates into Contribution Policy

The debate over actuarial discounting brings out a bit of schizophrenia over dual rates. It is increasingly (if grudgingly) recognized that the finance economists are right about discounting liabilities at a low-risk rate that corresponds to the guaranteed nature of promised benefits. And yet, the finance economists are typically careful to restrict their conclusion to reporting requirements, and not necessarily to funding policy.

Our analysis above points to an approach that at least informally integrates dual rates into a contribution policy: report liabilities λ accurately, using d, and then set a target funded ratio, f^*. As with standard actuarial policy, the asset accumulation goal is tied to liabilities, since that represents the cost of the benefits to which asset accumulation is directed. As we have seen from *(4)*, when the goal is set in this fashion, the required contribution rate is governed by both cost rates: the pay-go rate c^p and the (properly discounted) normal cost rate c^n.

The open question, then, is how to set the target funded ratio, f^*. For example, if we were to aim at reproducing current funding goals, represented by the target asset ratio of $a^* = 7$ depicted in Figure 4a, but with rediscounted liabilities, then the target funded ratio would be reduced from $f^* = 100\%$ (of wrongly discounted liabilities) to about $f^* = 60\%$ (of accurately discounted liabilities).

Our proposal, however, is to go back to fundamentals. In general terms, for public plans the target ratio should be based on the public's preferences for intergenerational cost-sharing and its tolerance for risk in pursuit of returns. We believe that our equation *(4''')* can be helpful in systematizing an approach to this decision, in conjunction with a semi-formal social welfare function. Let us posit the latter as $-V[(a^* - a_0), E(c^*), \sigma(c^*)]$, where $(a^* - a_0)$ is a short-hand measure of the costs required over some period to reach the asset target; $E(c^*)$ is the expected value of the steady-state contribution rate at which the asset target is aimed, given by *(4''')*; and $\sigma(c^*)$ is the risk associated with that target. Since these three arguments to V are social "bads," we preface V with a minus sign and let the partials V_1, V_2, and V_3 be positive.

The optimization problem over these three "bads" requires a joint decision on two instruments: (i) the investment allocation plan, formally represented by the target return r, and the associated risk premium $(r - d)$; and (ii) the target funded ratio, f^*. We do not propose here to spell out a full solution to this complex problem, but rather to sketch out the considerations that might generate such a joint decision and to infer some contours of what that would look like. Specifically, we consider the optimization of $-V[(a^* - a_0), E(c^*), \sigma(c^*)]$, subject to *(4''')*, over the choice variable f^*, conditional on the investment decision represented by $(r - d)$.

We first consider the polar case, where the plan has no tolerance for risk (V_3 is effectively infinite). In this case, the plan would invest entirely in fixed income, so r would be reduced to d, and the third term in *(4''')* would vanish. (Semi-)formally, the plan would only raise f^* so long as the marginal social benefit of a higher target exceeds the marginal social cost. Here, the benefit of raising f^* is the reduction in the steady-state contribution c^*, and the cost is the extra effort required to reach the target asset ratio:

Raise f^* as $-V_2\, dE(c^*)/df^* > V_1\, da^*/df^*$.

Using *(4''')* and $a^* = f^*\lambda^*$, we have:

Raise f^* as $V_2\,(c^p - c^n) > V_1\,\lambda^*$.

The key point here is that if liabilities are properly discounted at d, instead of r, so would their rate of accrual, the normal cost rate, c^n. This would raise c^n much closer to the pay-go rate c^p.[15] Consequently, in this case of *total* risk-aversion, there may be relatively little benefit $(V_2\,(c^p - c^n))$ to any marginal increase in the target asset ratio. The extent to which the target would be raised rests heavily on the degree to which future generations are weighed against current generations (V_2 vs. V_1). It seems unlikely that a totally risk-averse public plan would pursue full-

[15] The normal cost rate would actually exceed the pay-go rate if $d < g$. We assume $d > g$, but maybe not by much. Estimates in Lenney, Lutz, Schule, and Sheiner (2020) put rediscounted normal cost higher than the pay-go rate.

funding ($f^* = 1$), or anything approaching it, since that would require the accumulation of sufficient assets to generate fixed income flows matched to anticipated benefit payments. Undoubtedly, this would require massive hikes in contributions, imposing unacceptably high transition costs on the current generation for relatively little benefit in the steady-state.

However, public plans do appear to have some tolerance for risk by investing in assets with higher expected returns, r, but with greater risk. Our expression *(4''')* helps understand the additional considerations in play for optimizing f^*. Our expression for comparing marginal benefit and marginal cost would now be:

Raise f^* as $V_2\,[(c^p - c^n) + (r - d)\lambda^*] > V_1\,\lambda^* + V_3\,\rho(r - d)\lambda^*$,

where we take ρ as the standard deviation of return per unit of risk premium. Thus, in comparison with the polar case of total risk aversion, the additional marginal benefit is the extra expected return from a higher target asset ratio and the additional marginal cost is the extra risk. Clearly for any degree of risk-tolerance there is some range of f^* over which the benefit from the extra return exceeds the cost of risk. Thus, we would expect the optimal f^* to exceed that under total risk-aversion. Under what conditions (if any) we would expect the goal to be full-funding ($f^* = 1$) is not a question we can answer here. Our more modest intention is that the semi-formal expressions we provide can guide future empirical work to help integrate the insights of finance economics and steady-state analysis into pension funding policy.

Conclusion

Standard actuarial practice pursues intergenerational equity by employing funding rules that seek to ensure each generation pays for the services they receive. These rules do this through the concepts of normal cost and amortization, which together, in theory, should result in fully funded benefits for each cohort of workers and taxpayers. Normal cost is meant to pre-fund the full cost of benefits earned by a cohort of employees over their careers, while amortization is meant to close funding gaps that result from payment shortfalls and unrealized assumptions.

In practice, these rules have failed to adequately link earned benefits and contributions. The true market cost of earned benefits have been understated, leading to the accumulation of large pension debt and steeply rising contributions to amortize that debt. These payments are crowding out spending in other areas like infrastructure and education. Given this result, it is questionable whether current actuarial practice has effectively maintained intergenerational equity, as current generations are paying for past benefits. In addition, standard amortization practice often builds in a large drop in taxpayer contributions at the end of the amortization period. The current generation of taxpayers is arguably being asked to bear a disproportionate share of the atonement for past sins compared to future generations.

A primary cause of public pensions' current financial problems was the failure to adequately consider the risks involved and the implications of those risks and uncertainties for future generations of public workers and taxpayers. In this paper, we strive to better elucidate pension funding dynamics using basic parameters and steady-state analysis. We propose a new pension funding approach that allows for proper liability discounting, clearer consideration of risk, and smooth contribution adjustment. We believe our analysis may offer a more honest approach, both in properly discounting liabilities, and not promising a mortgage-burning party when contributions plummet. Finally, we sketch a social welfare framework that could be used to balance intergenerational equity, the quest for returns, and investment risk based on the sponsoring government's assessment of public preferences.

Our approach can be thought of as replacing both pieces of current actuarial practice: normal cost and amortization. Normal cost is effectively replaced by the steady-state contribution rate given in (4'''), which (i) allows for a blend between normal cost (properly discounted) and pay-go, depending on the target funded ratio; and (ii) allows for excess returns *(r – d)* in exchange for the risk borne by the sponsoring government.

The other part of our proposal effectively replaces amortization schedules, which are currently based on the quantity of pension debt, a set amortization period, and an assumed payroll growth rate to backload payments using the "percent of payroll" method. All of these elements have flaws, as the debt is understated by aggressive discounting, the amortization period sets a funding cliff, and the "percent of payroll" method often moves funding farther away from the target, through initial periods of negative amortization.

Instead of an amortization schedule, our approach sets out equation (2), which specifies the adjustment process to the steady-state contribution rate and target asset ratio. This ties to the first argument of the social welfare function sketched out above, spelling out the contribution trajectory required to reach any specified target asset ratio, such that the near-term burden can be weighed against the long-term (steady-state) reduction in contributions.

Our future work will delve into how we might operationalize the ideas we lay out here and the implications of real-world application. Specifically, we will explore the impact of stochastic investment returns on funding and contributions and implications of plan maturity and cash flow for our proposed funding approach. We hope that our pension funding analysis will help us learn from the sins of the past rather than repeating them in the present, imposing likely burdens on the future.

References

Costrell, Robert M., 2018. "The 80 Percent Pension Funding Target, High Assumed Returns, and Generational Inequity," *Contemporary Economic Policy*, Vol. 36, no. 3, pp. 493-504.

Costrell, Robert M. and Josh McGee, 2020. "Sins of the Past, Present, and Future: Alternative Pension Funding Policies," EdWorkingPaper: 20-272, Annenberg Institute at Brown University: August 7, 2020; conference paper for Brookings Municipal Finance Conference, July 13-14, 2020.

Lenney, Jamie, Byron Lutz, and Louise Sheiner, 2019a. "The Sustainability of State and Local Government Pensions: A Public Finance Approach," Brookings Municipal Finance Conference; 2019b; University of Chicago Policy Forum: the Pension Crisis.

Lenney, Jamie, Byron Lutz, Finn Schule, and Louise Sheiner, 2021. "The Sustainability of State and Local Pensions: A Public Finance Approach," Brookings Papers on Economic Activity, Spring, pp. 1-48.

Neusser, Klaus, 2021. Difference Equations for Economists.

Table 1: Pension Funding Notation
A = assets on hand
L = accrued liabilities, the present value of future benefits earned to date
f = funded ratio, A/L (full funding goal is f = 100%)
W = payroll
$a = A/W$ = assets/payroll
$\lambda = L/W$ = liabilities/payroll
***c* = contribution rate, % of payroll**
c^p = benefit payments as % of payroll ("pay-go rate")
c^n = newly accrued liabilities as % of payroll ("normal cost rate")
r = return on assets; $R = (1+r)$
d = discount rate used to calculate present value of liabilities; $D = (1+d)$
g = growth rate of payroll; $G = (1+g)$

Figure 1. Normal Cost, Contribution and Benefit Rates, FY01 – FY20
Public Plans Data: 119 state & 91 local plans

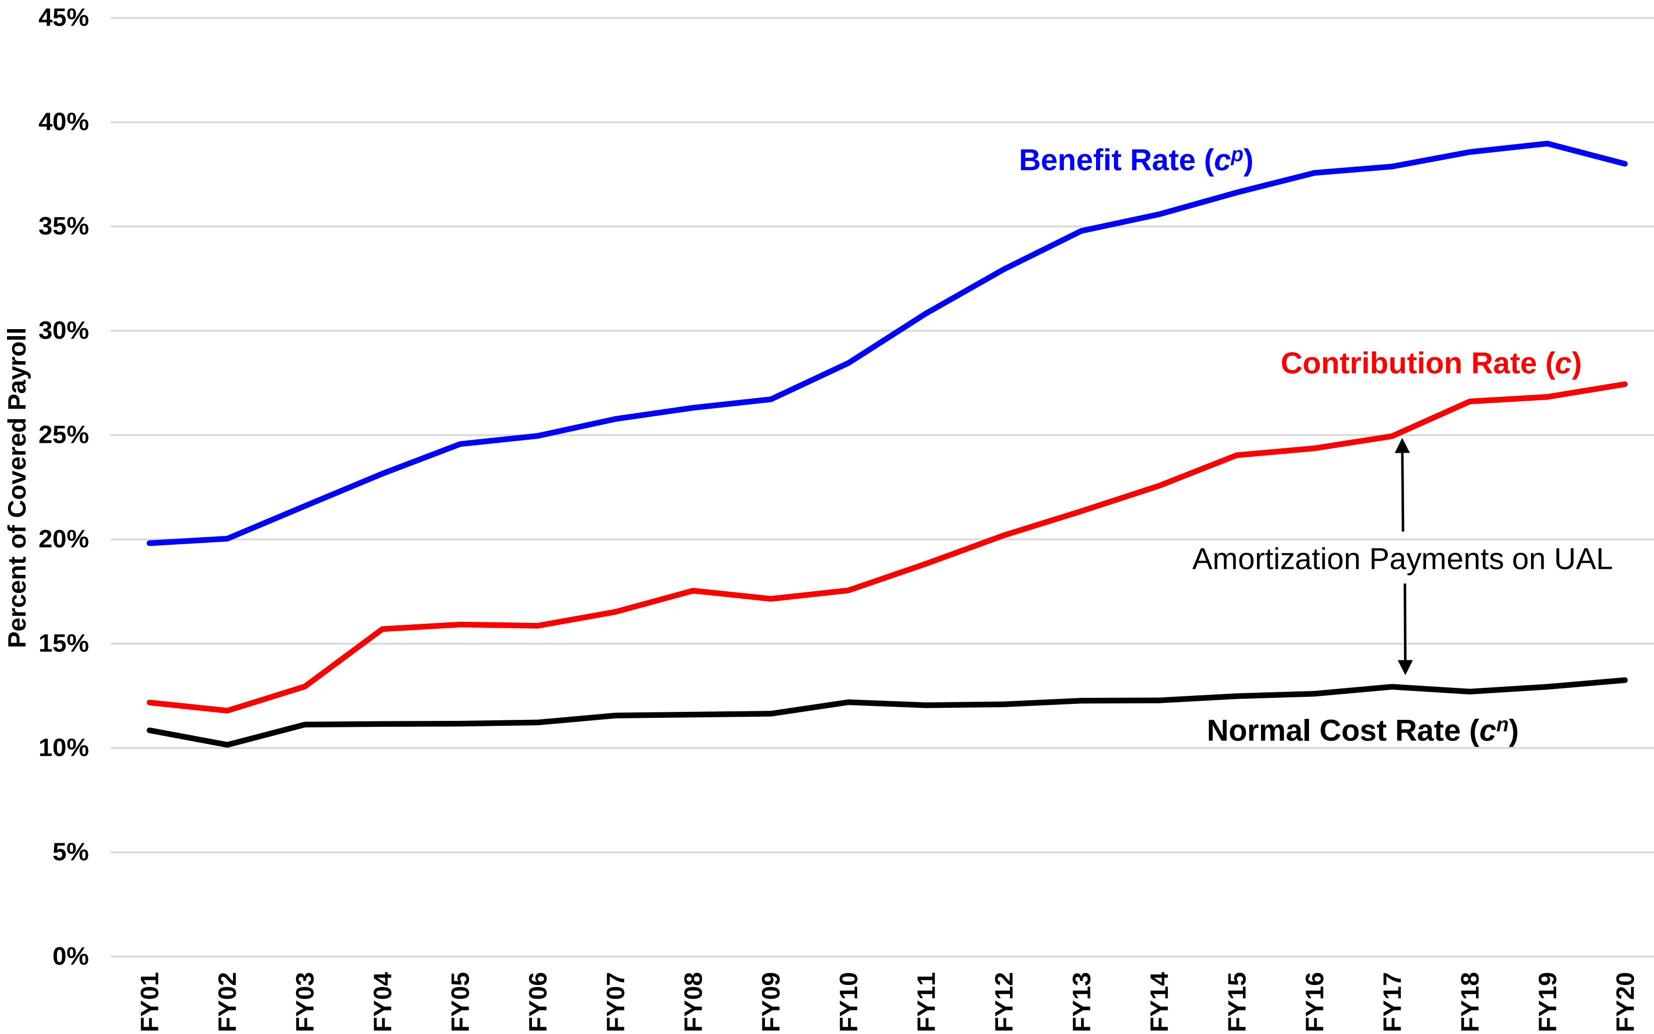

Source: Center for Retirement Research at Boston College
MissionSquare Research Institute, and National Association of State Retirement Administrators

Figure 2. Assets/Payroll, FY01 − FY20
Public Plans Data: 119 state & 91 local plans

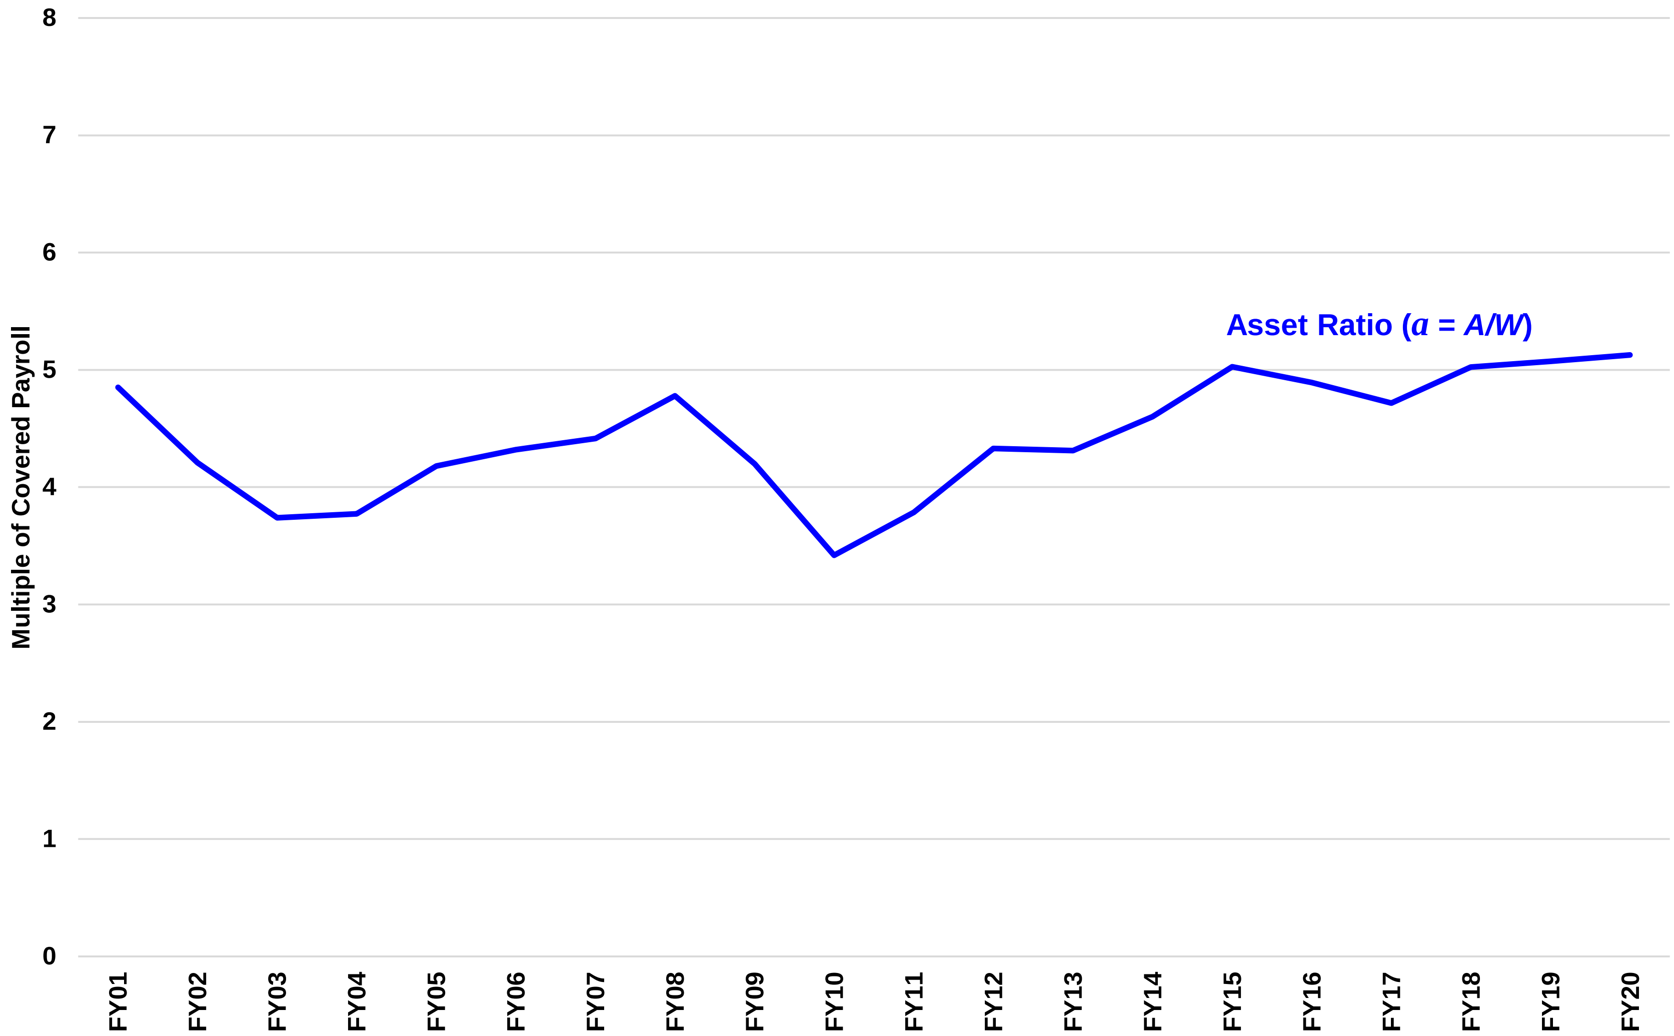

Source: Center for Retirement Research at Boston College
MissionSquare Research Institute, and National Association of State Retirement Administrators

Figure 3: Asymptotic Behavior of Asset Accumulation and Contribution Rate

r = 7%, g = 3%

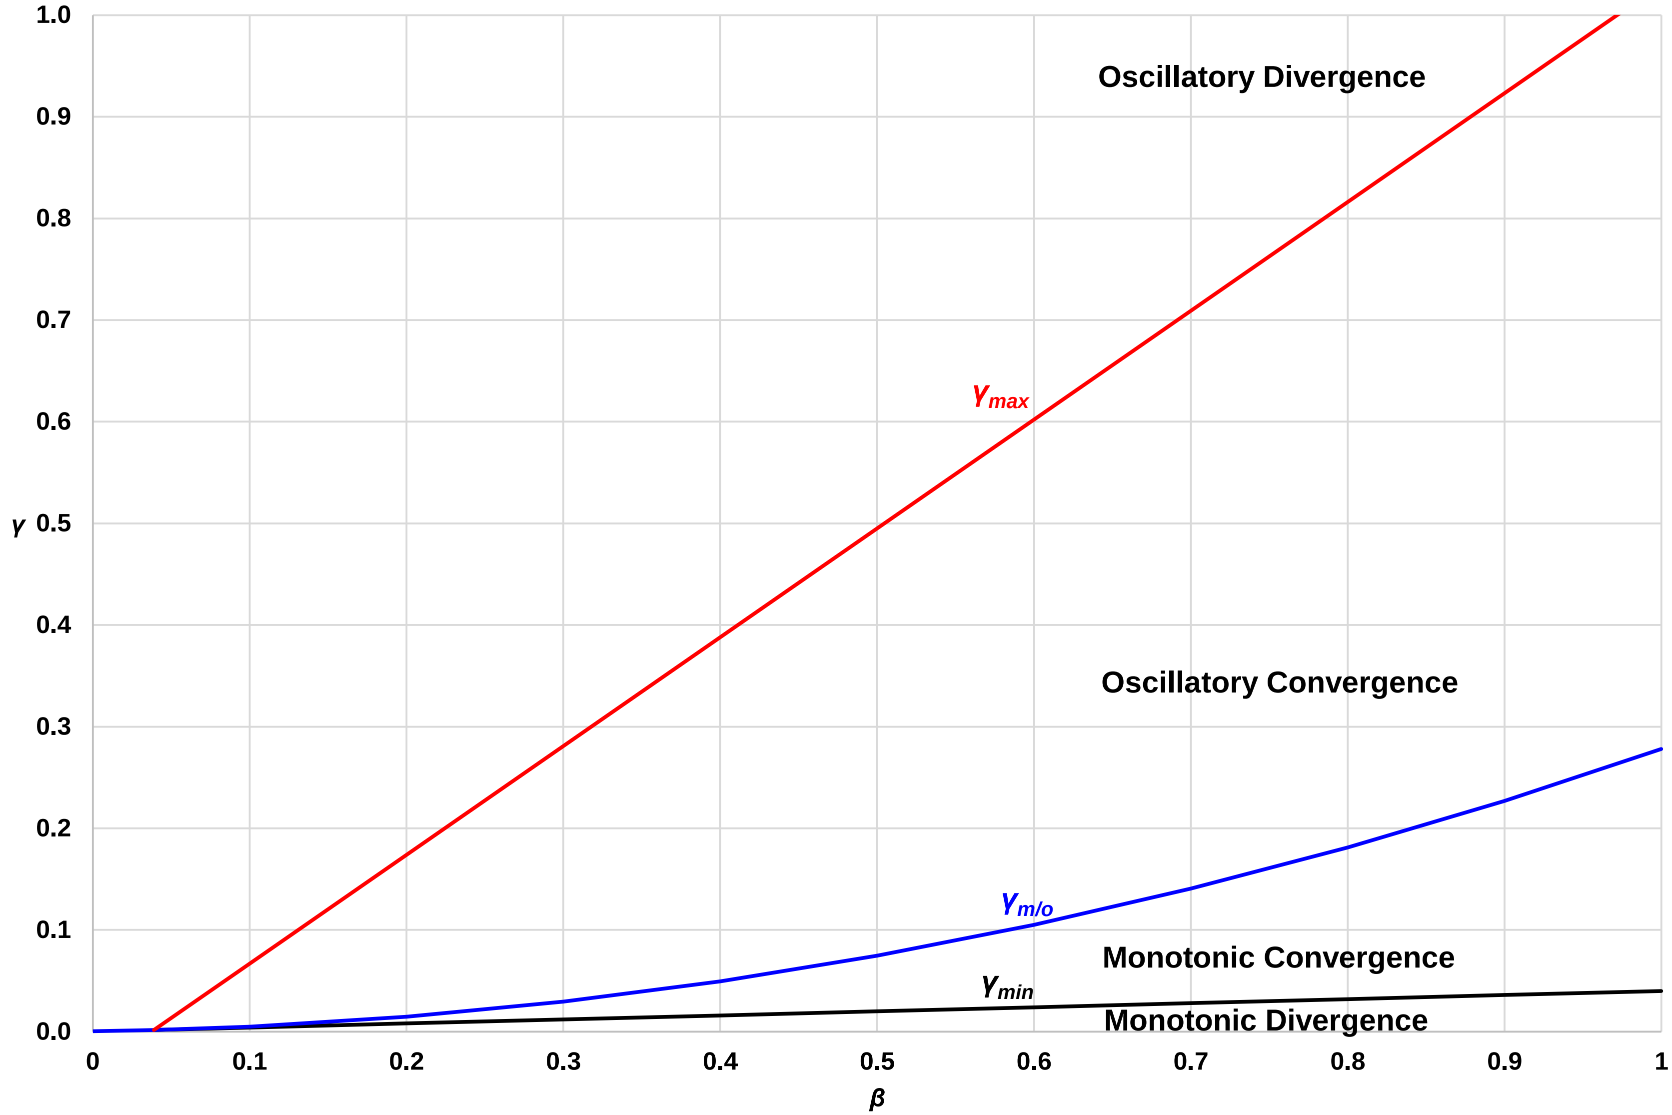

Figure 4a. Simulation of Contribution Rate & Asset Ratio

R = 1.07, G = 1.03, a^* = 7.0, c^* = 0.10, β = 0.5, $\gamma = \gamma_{m/o}$ = 0.075

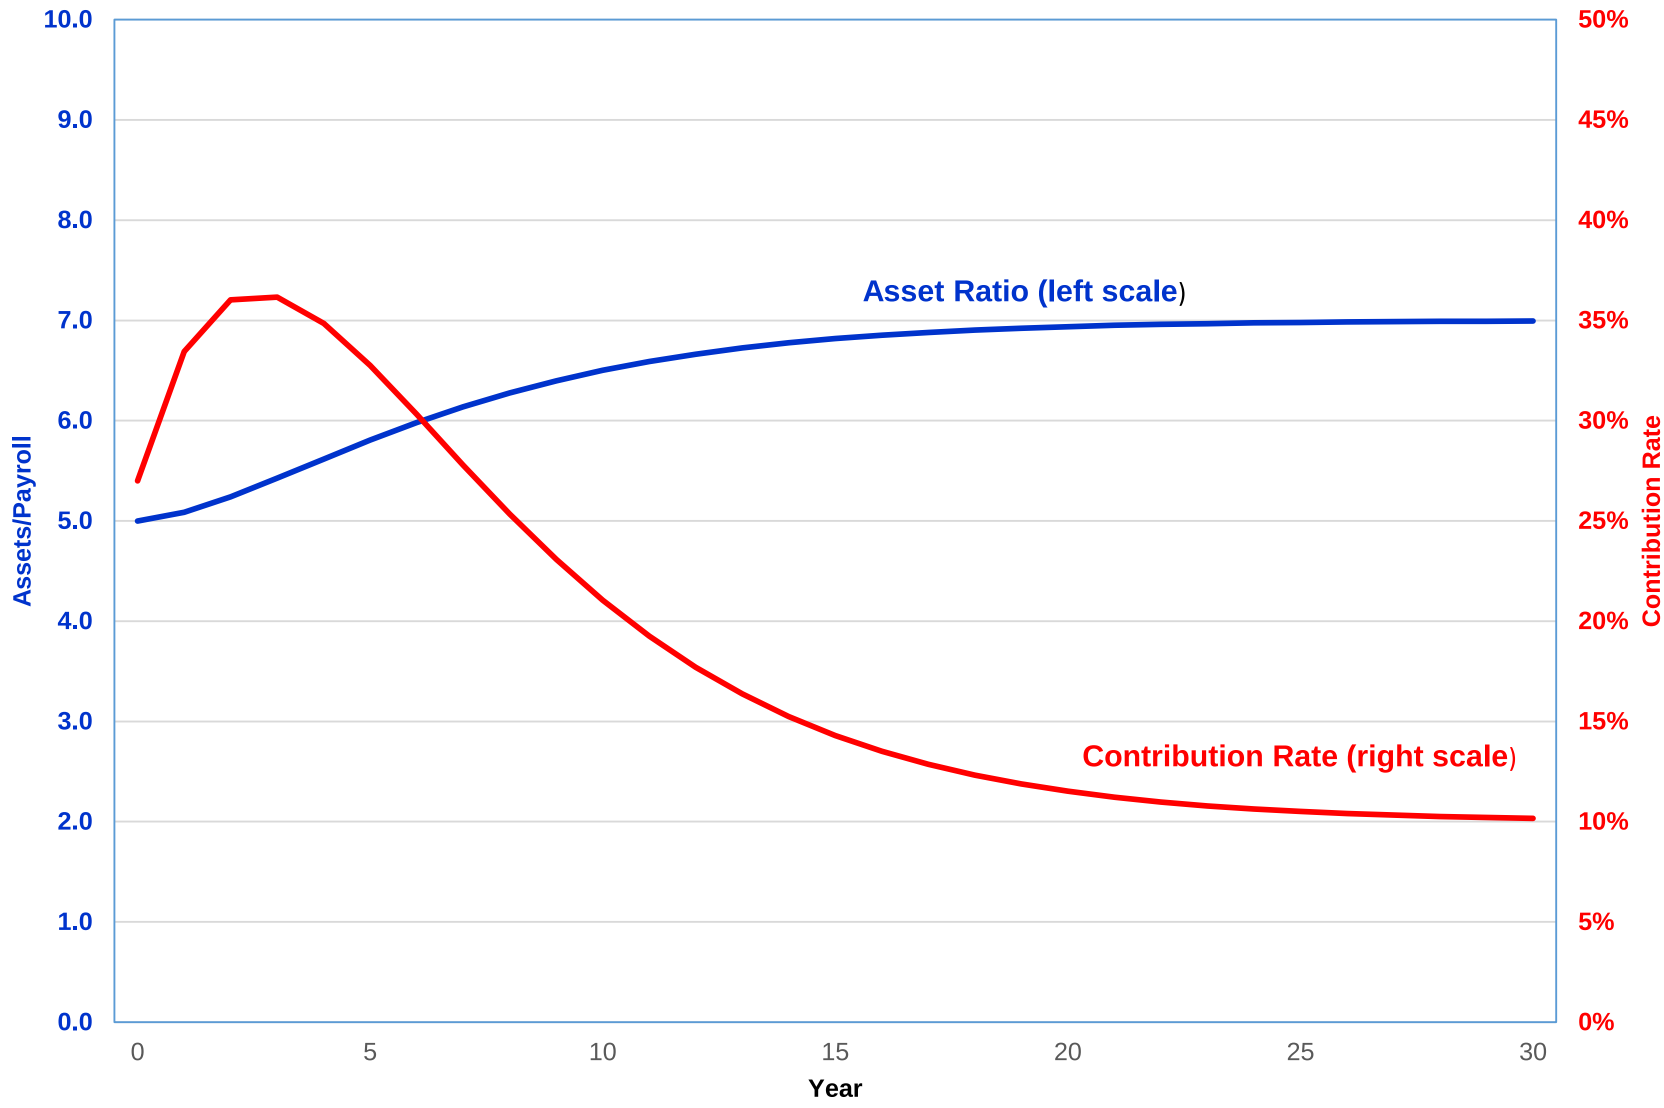

Figure 4b. Simulation of Contribution Rate & Asset Ratio

R = 1.05, G = 1.03, a^* = 7.0, c^* = 0.24, β = 0.5, $\gamma = \gamma_{m/o}$ = 0.069

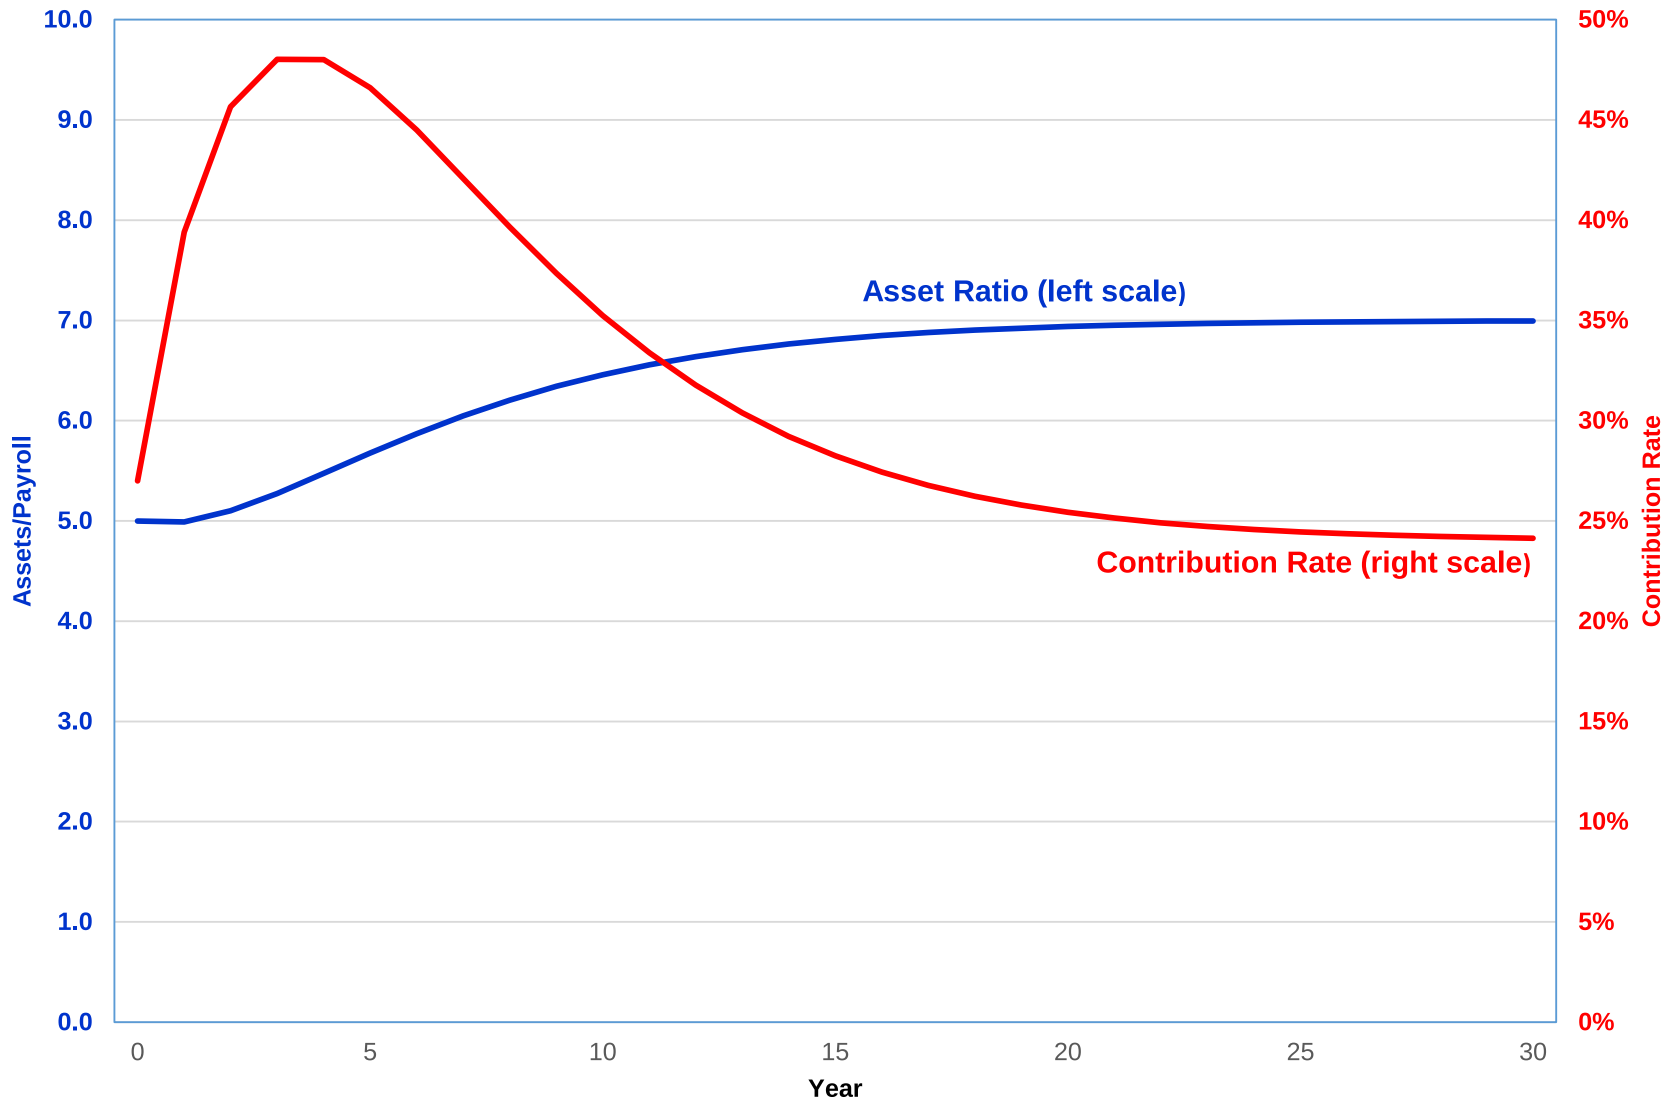

Figure 5. Assets & Liabilities, True & Reported, FY01 − FY20

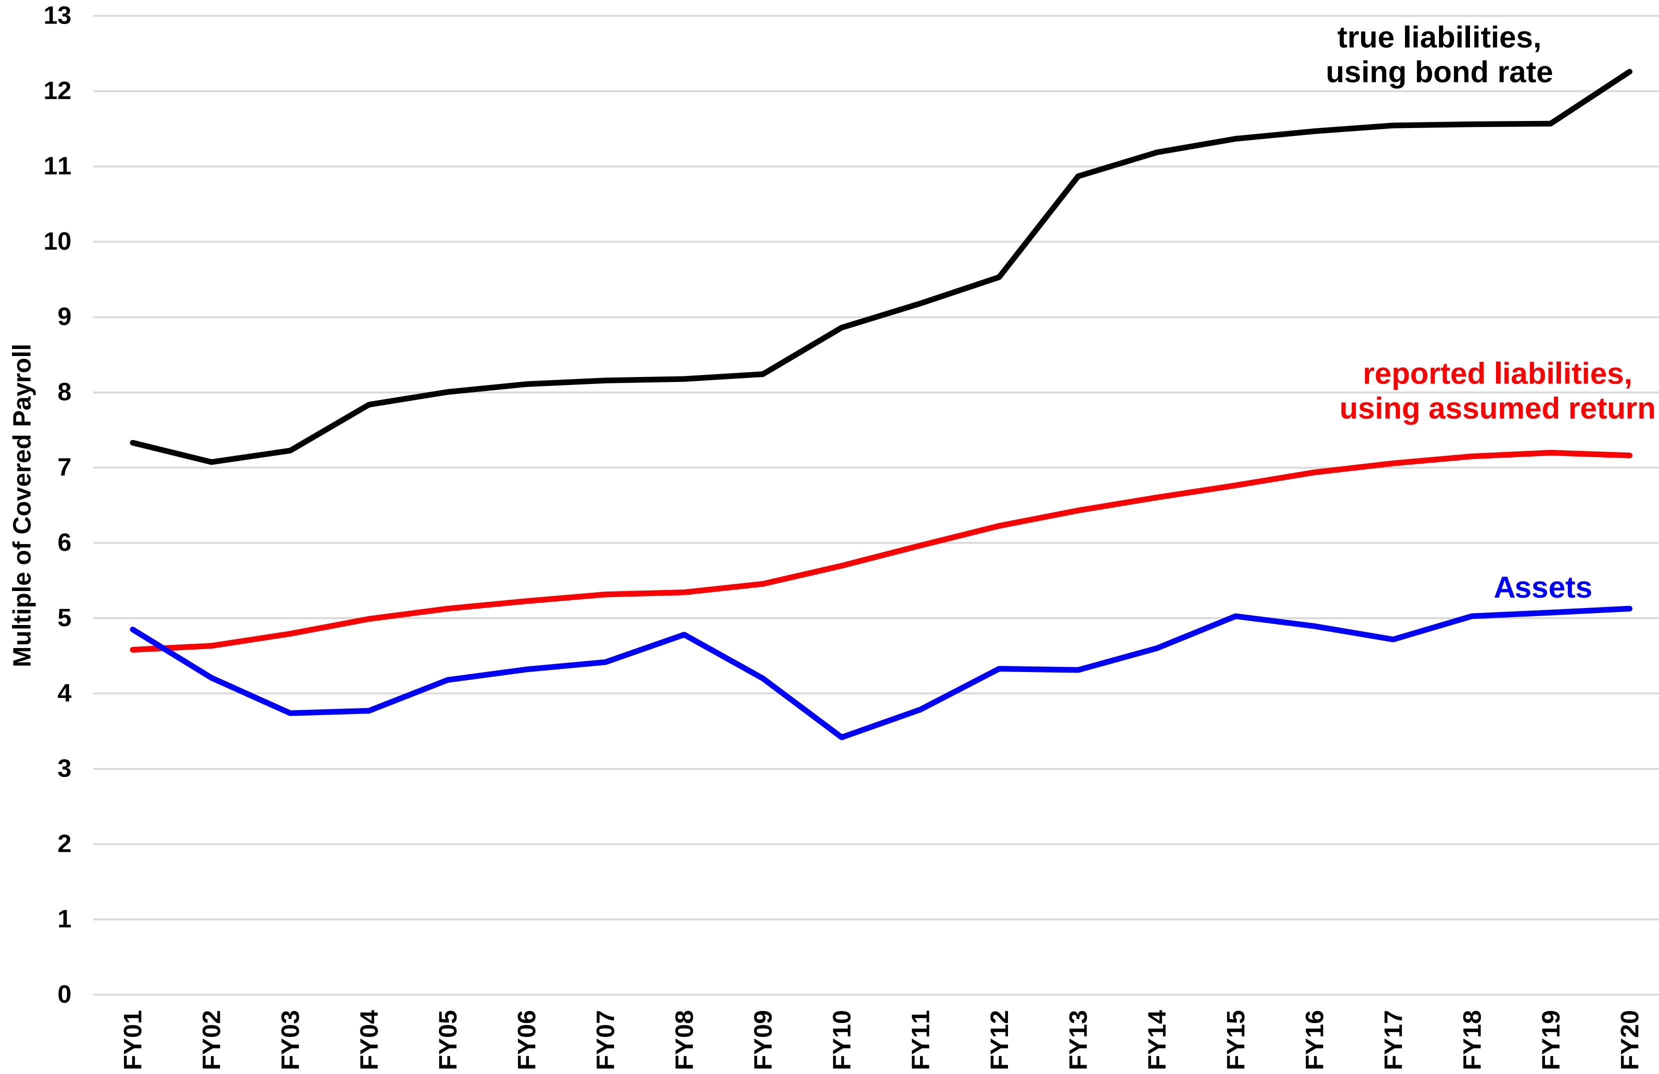

Sources: Center for Retirement Research at Boston College, Federal Reserve Board of Governors & authors' calculations
Both series use PPD payroll