Modeling Defined-Benefit Pension Plans: Basic Dynamics

By Robert T. McCrory, FSA
EFI Actuaries
1532 East McGraw St.
Seattle, WA 98112

Abstract

The purpose of this paper is to explore pension plan dynamics—the behavior of a defined-benefit pension plan over time in an uncertain environment. We proceed by simulating the behavior of a model plan in a stochastically varying economic environment. While both the model and its environment are simplified, the model displays interesting behavior suggesting policy considerations that should be included in the design and funding of these plans.

Keywords

Pension plans, actuarial funding, stochastic system dynamics, risk analysis

1. Introduction

A defined-benefit pension plan[1] is established by an employer to provide guaranteed lifetime benefits to retiring employees. Monthly benefit amounts are determined by a formula based on age and length of service at retirement and on the employee's pay. Disability and survivor benefits may also be provided. In establishing a defined-benefit pension plan, the employer is underwriting insurance: A benefit will be paid to qualified retirees—and possibly their survivors—for as long as they live, which could be several decades. The amount of the benefit is fixed, regardless of the economic or business climate faced by the employer.

Defined-benefit plans are important economically. State and local pension plans contain assets of over $2 trillion, provide benefits to over 7 million retirees, and generate $358.6 billion in economic activity, directly or indirectly producing 2.5 million jobs.[2]

Because the employer is issuing old age insurance, an actuary is required to determine a funding scheme whereby the promises inherent in the insurance can be delivered. Typically the actuary performs an annual actuarial study. Based on employee data the actuary computes the present value of promised benefits—liabilities—and compares those with the assets dedicated to fund them; the actuary then computes a funding scheme designed to bring assets in line with liabilities. This actuarial analysis is usually performed on a deterministic basis, with fixed assumptions used for rates of return on plan assets, inflation, pay increases, mortality rates and so on.

As noted above, the actuarial analysis is performed annually, providing a summary of the actuarial state of the pension plan once per year. Projections of plan behavior into the near future—five to 10 years—are becoming more common, but these projections are most often performed on a deterministic basis, with fixed assumptions. Little attention has been paid to the stochastic dynamics of pension plans. Specifically, how are pension plans likely to behave in a stochastic environment, with randomly fluctuating asset values and inflation? This is the subject to be addressed in this paper.

---

[1] Allen, E., J. Melone, J. Rosenbloom and D. Mahoney. 2002. *Pension Planning*. New York: McGraw-Hill/Irwin.

[2] National Institute on Retirement Security. 2011. "PENSIONOMICS: Measuring the Economic Impact of State and Local Pension Plans." National Institute on Retirement Security. [Online]. Available: www.nirsonline.org. [Accessed Dec. 28, 2011]

The dynamic behavior of a pension plan is most important in the area of governance. For example, we will see that there is enormous uncertainty in the amount required to pay for guaranteed benefits. Given the amount of uncertainty, a conservative approach to setting benefit levels is required, one that allows a significant margin for adverse variation. Regrettably, this has been lacking in many pension plan designs to date.

2. Methodology

In this paper we take an empirical approach: We will report on the results of a number of simulation experiments conducted on a Model Plan in a Model Economy. Both the Model Plan and the Model Economy will be highly simplified; nonetheless, they are representative of actual pension plans in operation, and we can learn a great deal from these simple models about the dynamics of actual pension plans in a real economy.

2.1 The Model Plan

To explore the basic dynamics of a defined-benefit pension plan, we will build a simple computer model of a pension plan. Below is a summary of its salient features.

- All members are the same.
  - All members are the same age when hired, and those who remain in employment retire when they reach a defined retirement age. In order to produce benefit levels and actuarial costs representative of a typical public sector plan, the hire age is 37 and retirement is at age 60.
  - All members receive the same salary beginning at hire, with annual increases for longevity and promotion of 2 percent for each year of service until retirement. In addition, members receive annual pay increases equal to the rate of inflation.
  - Members experience rates of mortality and termination that are typical in these pension plans.
- The member population is stable.
  - Each member who terminates employment, retires or dies before retirement is replaced by a new hire.
  - At time 0, the population has reached a steady state: The number of active and retired members at every age is stationary over time.
- Benefits are simple.
  - Upon retirement the member receives a benefit of 50 percent of final year’s pay as an annuity for life with a 2 percent annual cost of living adjustment (COLA). No survivor benefits are paid.
  - Deaths or terminations before retirement result in no benefit payments; there are no benefits other than the retirement annuity.
- All financial transactions—salary, benefit payments and employer contributions—occur annually at the beginning of each year. Investment earnings are paid once, at the end of each year.

The Model Plan is an open group retirement plan similar to those often found among public sector defined-benefit plans. Many details of actual plans—such as monthly benefit payments and ancillary benefits—have been omitted for simplicity, but that is unlikely to affect the overall behavior of the model.

This Model Plan is being funded in a fairly traditional way, using the Entry Age Normal Cost Method.[3] The actuary is assuming an annual return of 8 percent and an increase of 3.5 percent per year in the consumer price index (CPI) in her calculations. The plan begins at time 0 with a funded ratio of 100 percent: The market value of plan assets is exactly equal to the liability for benefits earned to date. All demographic assumptions (rates of mortality and termination, for example) are met exactly. The stochastic element of the model is limited to economics; the rest is

---

[3] Bowers, N., H. Gerber, J. Hickman, D. Jones and C. Nesbitt. 1986. "Theory of Pension Funding." Pages 533–555 in *Actuarial Mathematics*. Itasca, Illinois: Society of Actuaries.

deterministic.

Figure 1 below presents a simulation of the benefits paid from the Model Plan as a percentage of active member payroll over 100 years; 1,000 trials have been performed. Individual simulation trials are shown in gray, quartiles in blue, the mean in red, and the actuarial projection in green. In our simulation trials, all demographic assumptions are met exactly; the return and CPI assumptions are met on average by the Model Economy (see below). The actuarial projection in green is a forecast assuming all actuarial assumptions are met exactly; in particular, the return on assets is 8 percent and inflation is 3.5 percent in every future year, without exception.

Note that the range of benefit payments in Figure 1 is narrow: Values range from about 39 percent to 50 percent. Therefore, there is little variation in benefit payments as a percentage of payroll. The standard deviation of the benefit payments across all simulation trials and years is only about 1.7 percent. Note also that the distribution of benefits across all simulation trials and years is symmetrical, and the mean benefit payment is very close to the actuarial projection.

Variation as a percentage of payroll is observed because the COLA is fixed at 2 percent per year, while inflation may vary stochastically. Since benefits and payroll respond differently, benefits as a percentage of payroll are stochastic.

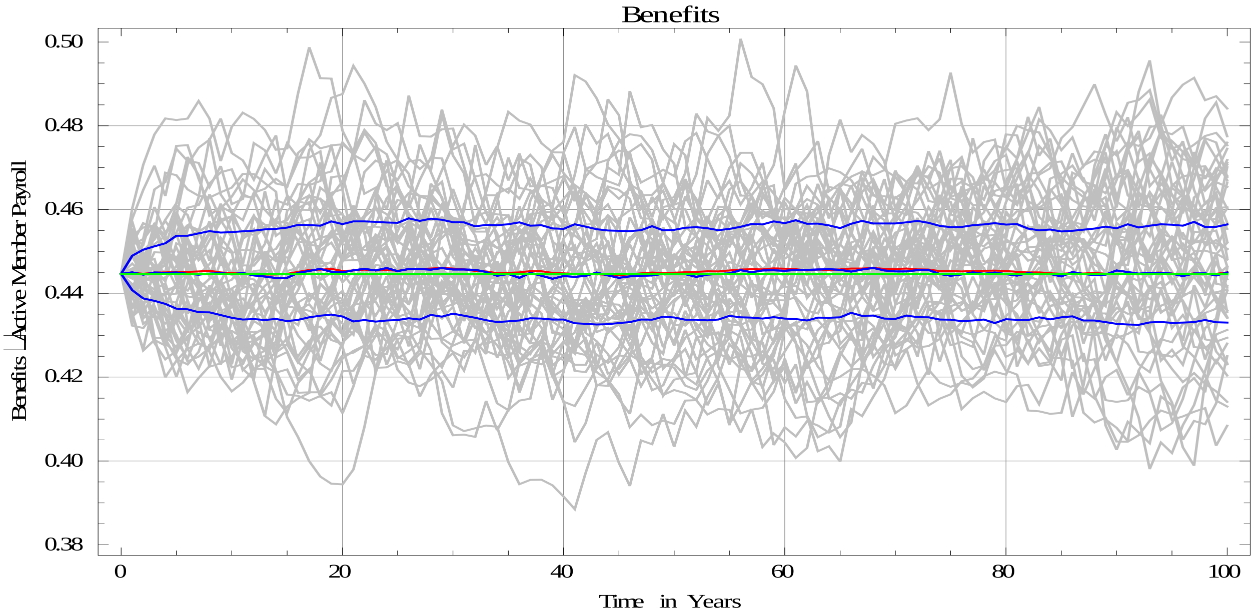

Year	0	25	50	75	100
Actuarial Projection	0.444688	0.444688	0.444688	0.444688	0.444688
Mean	0.444688	0.445407	0.444935	0.445368	0.445026
Standard Deviation	0.0	0.0167903	0.0166493	0.0166918	0.0170788
75th Percentile	0.444688	0.456838	0.455157	0.455785	0.456438
Median	0.444688	0.44532	0.444608	0.444202	0.444957
25th Percentile	0.444688	0.433621	0.434259	0.433709	0.433009

Figure 1: Simulated Benefits from the Model Plan as a Fraction of Active Member Payroll
(Each trial is shown as a gray line; the actuarial projection is in green; the mean is in red; and the median and quartiles are blue.)

2.2 The Model Economy

The plan operates in a stochastic financial environment, as follows:

- Annual investment returns are normally distributed with a geometric mean of 8 percent (8 percent compounded annually) and a standard deviation of 11 percent. This level of variability is typical of a 50 percent equity/50 percent fixed income mix as reported in the 2011 Ibbotson SBBI Yearbook.[4]

- Annual inflation is also normally distributed with a geometric mean of 3.5 percent (3.5 percent compounded annually) and a standard deviation of 1.5 percent. The standard deviation of the inflation rate is lower than the historical figure of 4.2 percent for 1926 through 2010,[5] but it gives a nice range of annual inflation rates that seem consistent with expectations that inflation will remain fairly docile.

- Inflation rates and investment returns are independently distributed.

This economic model is obviously highly simplified. Actual investment returns are not normally distributed, and how they are distributed is a matter of much study and debate. Recent experience has shown that extreme events are much more likely than would be assumed using a normal distribution. Moreover, the distribution and serial correlation of investment returns differ by asset type; consequently, these characteristics of the return will vary with different asset mixes. Therefore, different investment portfolios will behave differently.

To add to the unreality, this economy is everything the actuary could ever hope for. The actuarial assumptions for the Model Plan are met exactly, on average across all trials. However, for a single trial of 100 years, the geometric average return or CPI could differ substantially from the average across all trials. The Model Economy therefore represents a kind of actuarial wind tunnel—a fairly benign, controlled environment in which we can test the dynamic behavior of the Model Plan.

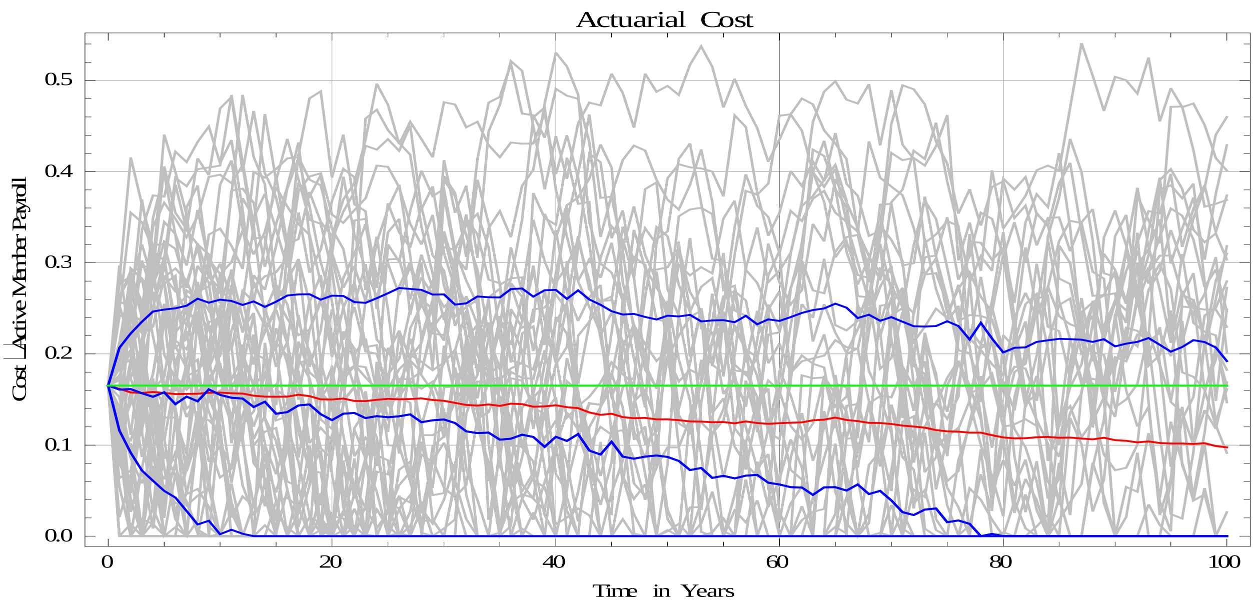

Year	0	25	50	75	100
Actuarial Projection	0.165275	0.165275	0.165275	0.165275	0.165275
Mean	0.165275	0.150708	0.128161	0.114998	0.0974617
Standard Deviation	0.0	0.143748	0.140221	0.139323	0.134221
75th Percentile	0.165275	0.266623	0.241916	0.235827	0.192109
Median	0.165275	0.13047	0.086934	0.015328	0.
25th Percentile	0.165275	0.	0.	0.	0.

Figure 2: Simulated Actuarial Cost of the Model Plan as a Fraction of Active Member Payroll
(Each trial is shown as a gray line; the actuarial projection is in green; the mean is in red; and the median and quartiles are blue.)

[4] Morningstar, Inc., Ibbotson SBBI 2011 Classic Yearbook, 2011, Chicago: Morningstar, Inc.

[5] Morningstar, Inc., Ibbotson SBBI 2011 Classic Yearbook, 2011, Chicago: Morningstar, Inc.

3. Cost

Figure 2 above shows the simulated actuarial cost of the Model Plan over 100 years. The actuarial cost is the amount determined by the plan actuary that must be contributed to the plan each year to pay for the retirement benefits promised by the employer. As in Figure 1, trials are shown in gray, quartiles in blue, the mean in red, and the actuarial projection in green. The investment return and annual inflation assumptions are met on average across all trials; however, in any single trial, the geometric mean investment return and annual inflation may differ substantially from the actuarial assumptions.

Most practicing actuaries would expect the average cost of the Model Plan to remain about level in the graph above: After all, our assumptions regarding the average investment return and inflation are being realized. However, something is pushing the average cost down over time: The average cost stays roughly level for 10 years or so, but then gradually decreases. The median cost begins declining sooner, and we reach a point at about 80 years at which more than half of the trials are producing no cost at all. Note that at any point in time, the 75th percentile is about twice the average cost.

This pattern of decreasing mean and median actuarial cost exists because of the effect of the Exclusive Benefit Rule.[6] The Exclusive Benefit Rule is part of the Internal Revenue Code governing qualified pension plans in the United States; it allows withdrawals from a pension plan only for the “exclusive benefit” of the plan’s members and beneficiaries. Refunds to employers in the event the plan becomes overfunded are forbidden.

The Exclusive Benefit Rule, when combined with actuarial funding methods, establishes an asymmetric contribution pattern: Contributions to the plan are required in times of poor returns and underfunding, but in times of good returns and overfunding assets cannot be withdrawn from the plan. The result is that overfunding is compounded rather than corrected during a series of favorable returns: Overfunding grows until the actuarial cost is zero, and is often compounded thereafter. On average this one-way street pulls the mean and median cost down over time.

4. Funding

Figure 3 below shows the simulated funded ratio for the Model Plan; once again, 1,000 trials have been performed. The funded ratio is the quotient of plan assets divided by the funding target, the actuarial accrued liability, set by the actuarial funding method.

Note the wide range of funded ratios in Figure 3: Many of the simulation trials soar well over 700 percent funding, some getting there quite quickly. As a result of such runaway overfunding, the mean funding ratio exceeds the top quartile of results within 50 years. Since funds cannot be withdrawn from the plan except in the form of benefits and expenses (the Exclusive Benefit Rule again), in trials with high investment returns the plan becomes overfunded and assets continue to grow, compounding without limit. The median funded ratio also increases, but is less affected than the mean by the extreme values that arise from overfunding.

Figure 4 below shows a distribution of the funded ratio at year 50. The funded ratio is strongly skewed to the right (skew = 1.53). This results from the asymmetric cost pattern enforced by the Exclusive Benefit Rule: Contributions are required to remedy underfunding, but contributions are not withdrawn when there is overfunding.

Note in Figures 3 and 4 that there is an effective floor of about 40 to 50 percent for the funded ratio. At this low level of funding, the required employer contribution is sufficient in itself to offset benefit payments plus any possible loss on the remaining plan assets. Under these circumstances, assets are bound to increase and fund insolvency is impossible—as long as the required actuarial contributions are made on time. A similar minimum funded ratio is observed in simulations of actual plans, so it is not an artifact of the Model Plan being used here.

Note also the very long tail in Figure 4: The funded ratios can equal or exceed 1,000 percent in some trials. As noted at the end of Section 3, the Exclusive Benefit Rule prevents withdrawals from plan assets other than to pay benefits. Once assets exceed about 10 times payroll, investment earnings are sufficient to pay benefits and add further to

---

[6] Internal Revenue Code, *Section 401(a).*

overfunding, even if there are no further contributions to the plan. In recent history, plan sponsors have sought to employ assets over the funding target to raise benefit levels which, as we will see later, increases the risk faced by the plan sponsor.

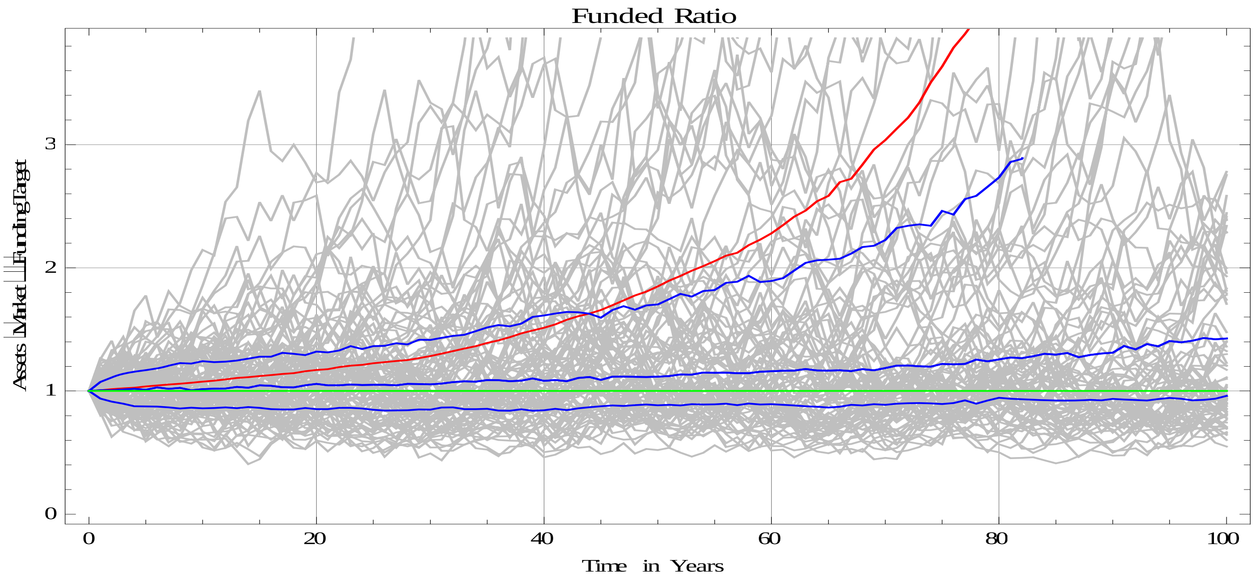

Year	0	25	50	75	100
Actuarial Projection	1.	1.	1.	1.	1.
Mean	1.	1.22554	1.84974	3.63337	9.48764
Standard Deviation	0.	0.648464	2.29238	8.29702	30.6395
75th Percentile	1.	1.36507	1.70268	2.46213	5.15995
Median	1.	1.05067	1.11554	1.22023	1.42662
25th Percentile	1.	0.846691	0.884112	0.894278	0.959494

Figure 3: Simulated Funded Ratios of the Model Plan
(Each trial is shown as a gray line; the actuarial projection is in green; the mean is in red; and the median and quartiles are blue.)

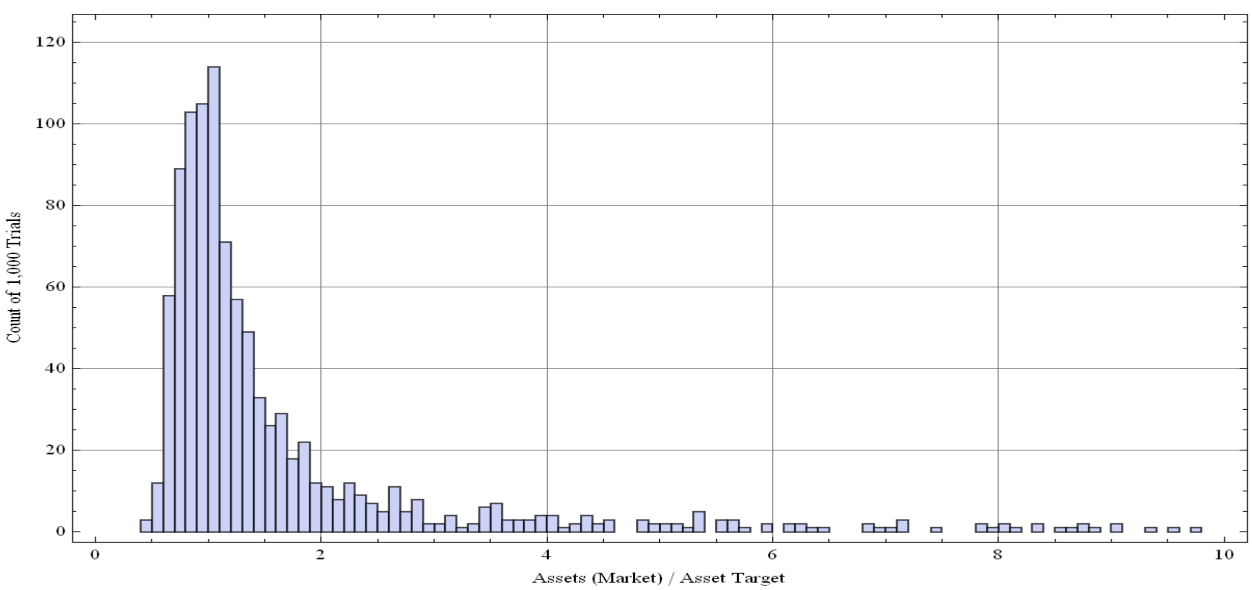

Figure 4: Distribution of Simulated Funded Ratios of the Model Plan at Year 50
(The vertical axis represents the number of 1,000 trials at the value on the horizontal axis.)

5. Risk

As we noted in the introduction, an employer sponsoring a defined-benefit pension plan relieves its employees of two major risks of retirement. First, while retirees with only a savings plan (a 401(k) plan or a tax-sheltered annuity) may outlive their savings, payments from a defined-benefit pension plan are guaranteed to continue until death. Second, while retirees with saving plans may experience wide variation in their fund balances with changes in the investment markets, a retiree's defined benefit continues unchanged. The employee has been relieved of the mortality and investment risks, and the employer has assumed them. Analyzing the dynamics of a defined-benefit plan includes identifying and quantifying the risks assumed by the employer.

The mortality risk mentioned above is usually not immediate. A defined-benefit plan is a risk pool: Gains from retirees who die when relatively young offset losses from their longer lived colleagues. Therefore, the mortality risk arises gradually from a long-term overestimate of the number of deaths after retirement, not from any sudden event. However, there are a number of other risks that can emerge quickly and unexpectedly, stemming from the investment of plan assets and the method used to fund a defined-benefit plan. In this paper we will consider two of these risks.

- Investment Risk

  Pension plan assets can be invested in virtually any kind of assets, ranging from government bonds to timberland. A pension plan is invested in riskier assets in the hope and expectation that a reward will follow in the form of higher returns and lower pension costs. In Section 5.1 below we will present a simplified market risk and reward trade-off to assess the impact of asset selection on plan risk.

- Funding Risk

  As the plan's assets increase, its funding status improves, and assets increase relative to payroll. While this is viewed as a favorable development, changes in the higher asset value now represent a larger percentage of member payroll, with a concomitant increase in the impact on cost as a percentage of pay. In Section 5.2 below we will look at the role of improving funding on increased risk.

5.1 Investment Risk

As noted above, in sponsoring a defined-benefit pension plan, the employer assumes an investment risk: No matter how much or how little is earned on the assets invested on behalf of the pension plan, the benefits paid must continue unchanged. Consequently, plan sponsors are confronted with the classic risk/return dilemma: Riskier assets, such as stocks or real estate, provide the possibility of an additional return, offsetting the cost of benefits, but they do so at the expense of unpredictable variation in asset values and actuarial costs. An employer investing in safer assets will enjoy a predictable, but low, return on assets, resulting in higher, but more predictable, employer contributions to the pension plan.

It would be useful to quantify—even if only approximately—the risk/reward trade-off characteristic of the asset markets in which pension plans invest. To this end, we have taken the capital market assumptions from a well-known investment consulting firm and we have simulated the performance of a wide range of asset portfolios of varying risk. The result is shown as the blue points in Figure 5 below.

Figure 5 is the risk/reward graph plotted for this set of capital market assumptions. Geometric average return for a 10-year period is plotted on the vertical axis with annual standard deviation of return on the horizontal axis. There are 10 asset classes in the model; we are provided with arithmetic and geometric mean returns, standard deviations of annual returns, and an annual correlation matrix. Total portfolio return is modeled as a joint normal distribution.

To determine the risk/reward trade-off, we find all allocations among the 10 asset classes, where steps in the allocation to each class are in 25 percent increments. A count gives us 715 portfolios in which the individual asset class allocations sum to 100 percent. We simulate the returns from each allocation for 10 years, using 250 trials. We then plot the 10-year geometric average return against the standard deviation of all simulated annual portfolio returns. This produces Figure 5.

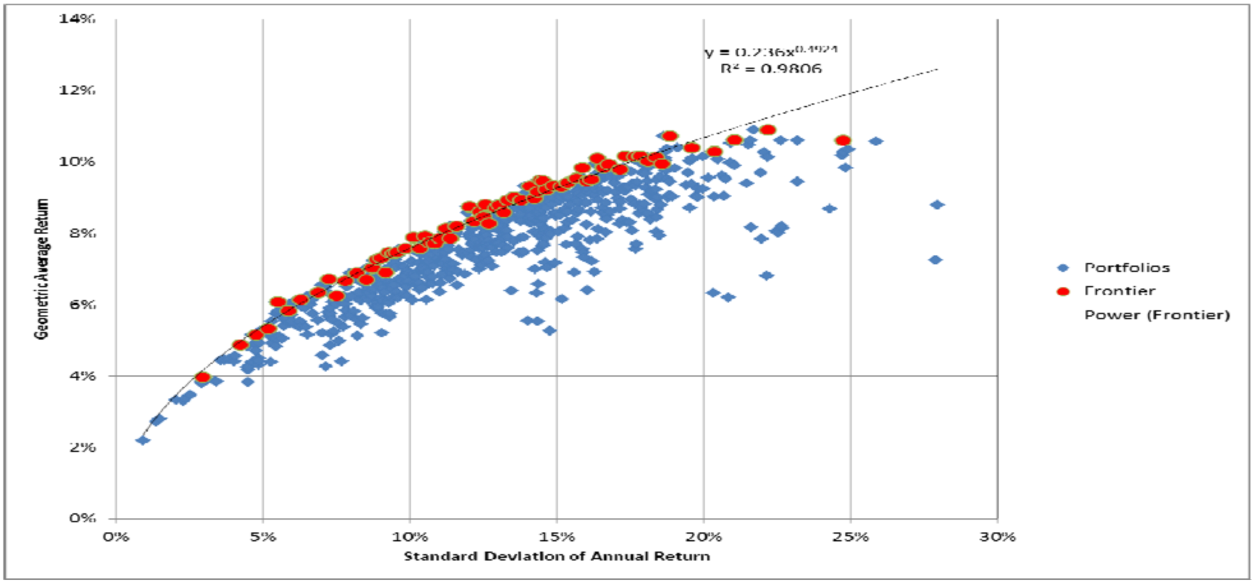

Figure 5: 10-year Geometric Average Return Plotted Against Standard Deviation of Return for 715 Portfolios (The horizontal axis is risk (standard deviation of return) and the vertical axis is reward (mean return). Points in red show the efficient frontier and a power function approximation.)

In studying Figure 5, we observe the traditional convex-down shape and the diminishing gains in returns with riskier portfolios. We also note that the upper left edge of Figure 5—the efficient frontier—passes very close to an 8 percent compound return with an 11 percent standard deviation, which has been our assumption for the Model Economy so far. The efficient frontier is significant in that it consists of a series of investment strategies that offer the highest return for a given level of risk, as measured by the standard deviation of return.

It would be useful to have a simple equation that relates the expected compound return to the risk. To this end, we approximate the efficient frontier. To do so, we sort the 715 portfolios by risk, divide the sorted portfolios into groups of 10, and plot the highest return in each group against the largest risk. These are shown in the red points in Figure 5. Then, we fit an equation to the plotted series of points.

A power function with Return = 0.236 × Risk $^{0.4924}$ does a very good job of describing the functional relationship of the frontier points, and we will use this equation as a measure of the risk/return trade-off. It is interesting that the return produced by this equation for an 11 percent standard deviation is 7.96 percent. I should point out that this agreement is coincidental: The capital market model is from an unrelated firm, and it was not developed with any input from me or my firm.

Figure 6 plots the mean and standard deviation of the actuarial cost versus the risk level as measured by the standard deviation of the return of the invested plan assets. Note that average cost decreases with increasing investment risk, but that the rate of decrease slows as the risk grows larger. Note that the standard deviation of cost increases with increasing investment risk, but not after 10 percent or so, after which the standard deviation of cost remains constant even as investment risk doubles. Note that the mean and standard deviation of cost cross at about 11 percent.

Figure 6 could be referred to as the cost risk/reward profile of the plan: It displays the reduction in actuarial cost and the increase in cost variability that comes with investing in riskier assets. Note that once the return risk increases above 10 percent or so, there is little additional performance—in terms of lower cost or a change in the variability of cost—as you take more risk. Moreover, in Figure 5, we notice that at the higher levels of risk—over 20 percent—actual modeled portfolios have lower mean returns than the power function fitted to the efficient frontier would predict. Therefore, at higher risk levels plan performance may be worse than that shown in Figure 6.

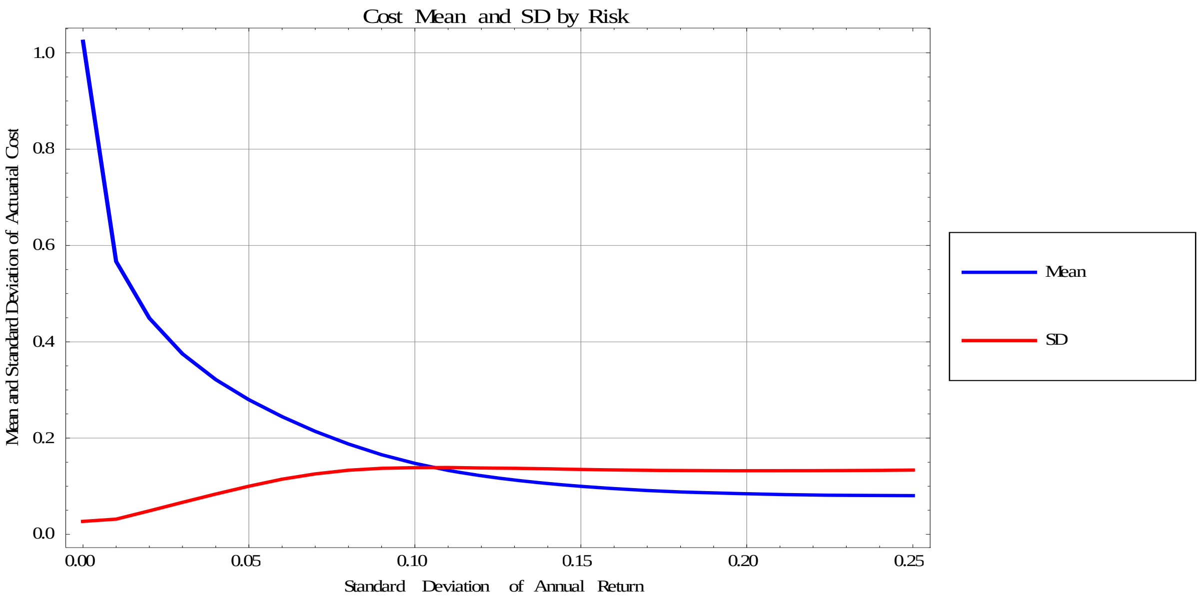

Figure 6: Mean and Standard Deviation of the Actuarial Cost Plotted Against Asset Risk

5.2 Funding Risk

Any change to a defined-benefit pension plan that increases assets relative to active member payroll increases cost volatility. In the case of higher benefits, the increase in benefits raises the actuarial funding requirement, the funding target used to determine costs. The resulting costs push assets up to meet the higher funding target, resulting in a higher ratio of assets to payroll. As the ratio of assets to payroll increases—this is called the "stability ratio" in some systems—investment gains or losses are a larger percentage of covered payroll. Accordingly, plan cost as a percentage of payroll is more sensitive to market variations.

The key point to note is that *any* increase in assets relative to payroll increases cost volatility. An unfunded plan has no risk from market fluctuations; an overfunded plan is exposed to much more risk. Let's see how the standard deviation varies by the funding status of the plan. Specifically, does cost become more or less uncertain as the plan's funded ratio increases?

Figure 7 shows cost standard deviation as a function of the funded ratio. This graph was constructed by taking the simulation results from the Model Plan for cost and funded ratio (Figures 2 and 3 above) and pairing each observed funded ratio with the following year's actuarial cost. With 1,000 trials of 100 years each, there were 100,000 pairs of funded ratio and succeeding year's cost. These points were sorted by the funded ratio and partitioned into 100 sets of 1,000 points each; for each set of points, the standard deviation of next year's cost was plotted against the mean of the funded ratios of the points in that set.

Figure 7 quantifies the observation that increasing assets leads to increasing cost volatility. We note that as the funded ratio increases from 50 percent to about 105 percent, the standard deviation of next year's plan cost nearly doubles from around 4 percent of payroll to over 7 percent. This variability should be compared with the mean cost of around 16 percent of pay for our Model Plan. After reaching a peak at about 105 percent funded, the cost standard deviation declines. As the plan becomes better funded, with more assets, and the actuarial cost declines, variability is reduced by the impossibility of a negative cost, due to the Exclusive Benefit Rule.

Note the standard deviation here is lower than we observed in Figure 3. Previously we have been considering a global standard deviation, independent of funded ratio. Here we are considering a conditional standard deviation of cost, in which the funded ratio is constrained to be within a certain range of values. Naturally, the standard deviation is lower in the constrained-conditional case. The important point here is that improved funding is not an unalloyed blessing: It brings with it increased cost sensitivity to investment fluctuations. The more you have, the more you can lose.

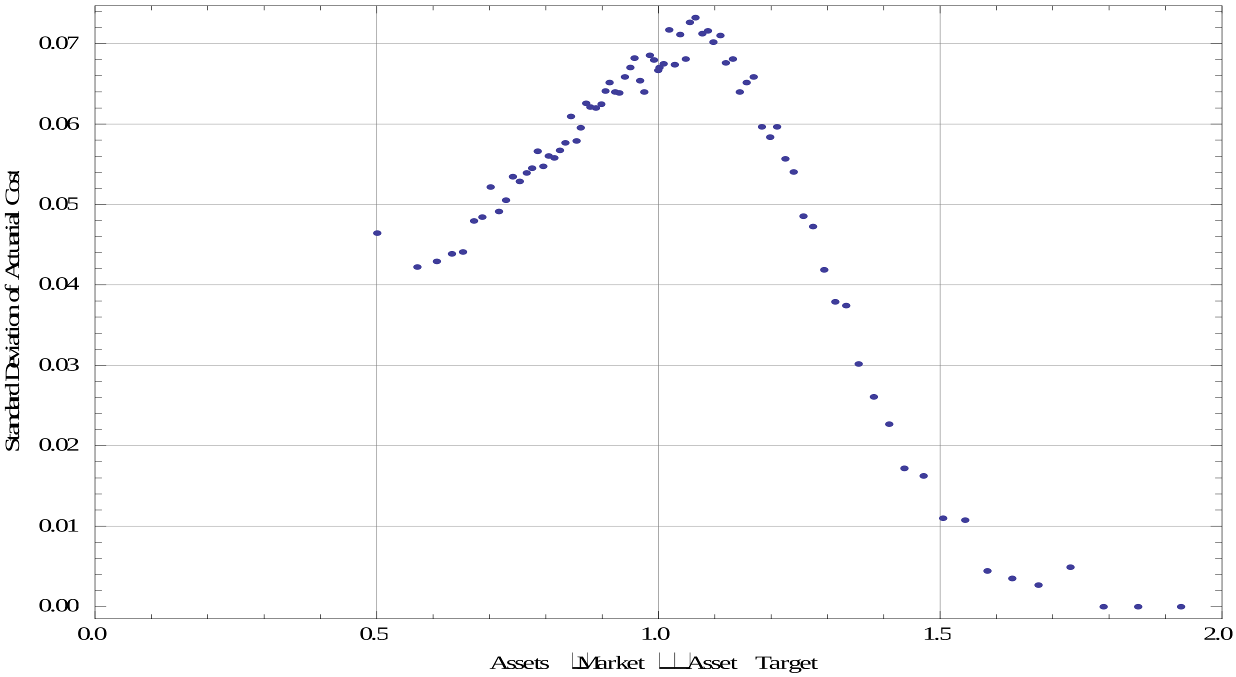

Figure 7: Cost Standard Deviation Plotted Against Funded Ratio

6. Summary and Conclusions

Defined-benefit pension plans are complex dynamic systems, and they display complicated and counterintuitive behavior. To summarize what we have discussed:

- Sections 1 and 2 introduced the simplified Model Plan and Model Economy used in this paper.
- In Section 3 we found that the actuarial cost of the Model Plan is neither level nor stable, even when the actuary is right overall (the actuarial assumed return equals the geometric average return over all trials).
- In Section 4 we observed that the level of funding of a pension plan can vary over a very wide range, but has a lower limit at which funding must improve, as long as the employer makes the required actuarial contributions.
- In Section 5 we discussed a couple of the risks faced by a pension plan and its employer sponsor. We discovered that riskier investments bring more cost variability, and that cost variability increases to a maximum when the plan is roughly 100 percent funded.

Understanding the dynamics of defined-benefit plan is crucial to their governance.

- In establishing the benefit promised by a defined-benefit plan, the likely variability of plan cost shown in Figure 2 must be carefully considered. A deterministic actuarial cost projection is not enough; a margin must be maintained between the average annual cost computed by the actuary and the amount the plan sponsor can afford to pay. This margin is essential to allow for the variation in cost that will occur due to the uncertainty of investment returns.
- In the past, when assets have increases and the funded ratio of the plan has exceeded 100 percent, plan members and employers began discussing how to use the “excess” assets over 100 percent to increase benefits. As Figure 7 shows, it is at this point that the cost of the plan displays maximum uncertainty, and today’s surplus may be eliminated by the investment markets overnight, as happened in 2000 and 2008. It is when the assets are highest that risk is highest, and great caution must be exercised.

- The risk/reward trade-off in investing plan assets is, with the benefit level, one of the two most important decisions to be made. One can adopt a low-risk/low-return/high-cost approach by investing in bonds, or a high-risk/high-return/low but variable cost approach using other assets. Understanding the implications of this decision shown in Figure 6 can be helpful is assessing the risks and rewards. In addition, Figure 6 and 7 taken together may suggest that as the plan assets increase and the funded ratio rises, the asset mix may move from risky assets to a portfolio with less risk, designed to offset the higher cost variability that accompanies higher funding.

Pension plan design and governance must take place in the context of the dynamics of the plan.