Pension Funding Without Liabilities

Robert T. McCrory, FSA, MAAA, FCA, EA

Presented at The Great Controversy: Current Pension Actuarial Practice in Light of Financial Economics Symposium
Sponsored by the Society of Actuaries

Vancouver

June 2003

(note: The quotations appearing in this monograph are exact, except where capitalization and punctuation were changed in keeping with modern style and grammar guidelines.)

Abstract

Determination of annual pension contribution has traditionally been accomplished via the computation of the "liabilities" of the plan. Such liabilities are the discounted present values of all or a portion of projected plan benefits. Computation of discounted present value poses no difficulties when the financial environment in which the plan operates allows selection of a reasonable discount rate. However, in environments in which returns on plan assets are highly variable, selection of a discount rate is difficult. Moreover, continuing to use discounted present values in a variable financial environment may be an inappropriate use of the tool. It *is* possible to compute annual pension costs without computing liabilities first, or at all, even implicitly. Moreover, such an approach can be shown to have powerful advantages in stochastic financial environments.

1. Introduction

My goals for this paper are quite modest:

- By using an example, I will show how a pension plan can be modeled on a realistic basis. The model will include two vital aspects of defined benefit (DB) pension plans that are often omitted from our actuarial models: Future new entrants and stochastic variability of asset returns and inflation.

- I will develop an approach to funding the example plan. In doing so, we will note that present value and liability calculations do not arise naturally and are not needed.

- I will urge that DB pension plans be approached empirically—with an attitude of exploration and experimentation—rather than analytically with simplified mathematical models. The mathematical models we typically use are too abstract and leave out too much.

This paper is not meant to be an academic exercise. Instead, it is written for working actuaries who are responsible for modeling pension plans as part of their livelihood. I hope this paper will suggest ideas and techniques that will be useful and accessible, so that actuaries can apply them in their practice on a regular basis.

Accordingly, I have written this in a style that I hope you will find approachable and even friendly. But be careful: I may have slipped up and included some real content here that will reward your consideration and thought. Let's get started.

2. An Example

We begin with a large statewide DB plan. As of July 1, 2000, some of the principal statistics for this plan are shown in Table 1.

Table 1

Example Plan Characteristics

Number of Active Members	254,124
Covered Active Payroll	$7.1 billion
Annual Benefits Paid	$1.2 billion
Assets	$31.6 billion
Equity/Fixed Allocation	About 70%/30%
Funding Method	Entry age normal
Return/CPI Assumptions	8.25%/3.5%
Benefit Formula	2% of highest year pay at age 55: higher percentage at higher retirement ages, lower percentage at lower retirement ages
Cost of Living Adjustment	Inflation, but not to exceed 2%

Now let's suppose we don't know anything about being an actuary (which is probably true in my case). Our mission (should we decide to accept it) is to determine some sensible way to pay for the plan benefits. How should we proceed?

2.1 Project Plan Benefits

As a first step, it seems natural to project the future benefits from the plan. If the plan is expected to continue indefinitely—as this statewide plan is—then the projection should be for a fairly long period. A long projection will mean that the benefit payments to future plan members must be considered. Accordingly, an open-group projection—assuming future new entrants—is appropriate.

(In practice, I have found that a projection of at least 50 years is desirable; I usually use 100 years. During the first 30 years or so, the projection of future benefits is strongly influenced by the actual demographic mix of the plan in question. After that, projected benefits and payroll are increasingly determined by the actuarial assumptions used in the projection. Therefore, the projection after 30 years can be used to reveal inconsistencies in the actuarial assumptions. Even modest inconsistencies can cause runaway growth or decline in projected benefits or payroll.)

In the event that the plan is known or expected to terminate or be closed to future new members at some future date, this fact should be included in our projections.

So, as a first step, we need to develop a model of our plan.

2.2 Develop a Plan Model

Our plan model will be very similar to the models currently used by pension actuaries. Even starting from scratch, it is not difficult to appreciate that our projection of benefits will require observation of past patterns of retirement, termination, disability and mortality, and the projection of these patterns—perhaps with some modification—into the future.

Patterns of pay increase due to longevity and promotion can be observed and projected. Future active workforce levels also need to be projected, usually with the help of the plan sponsor. If there are multiple tiers of benefits—with new members receiving different benefits than current members—then terminations among older tier members will cause new entrants in the current tiers. This shift in the workforce from older to newer tiers should be part of our model in order to reflect the changing level of benefits over time in our projections.

But we cannot proceed without including some projection of future inflation. Inflation will drive the growth of active payroll, and it may also affect benefits through the mechanism of a cost of living adjustment (COLA). This means that we are forced to posit some type of economic model. So we have to take a step back in order to move forward.

2.3 Develop an Economic Model

An economic model is essential in projecting the financial future of any pension plan. For most of our work as pension actuaries, we use a very simple economic model: A single rate of return on plan assets and a single rate of inflation.

In general, the models used by actuaries can be—and probably should be—fairly simple. Among the characteristics I would like to achieve in an economic model are the following:

- *Comprehensive* (enough). The model should include projections of those economic variables that are necessary for our projections. Inflation is mandatory, as is the overall investment return on plan assets. Separate investment returns for major asset classes can be useful in measuring the effect on the plan of different asset allocations. Interest rates at various durations—such as the notorious 30-year Treasury bond or its surrogate du jour—may be useful, especially in private sector plans. Recently, I have received requests from plan sponsors and boards to include investment income in the form of interest and dividend payments; they want to estimate when the outflows from plan funds (benefits and expenses) are likely to exceed inflows (contributions and investment income). Depending on the circumstances, other economic variables may also be included. For example, one of my plans is funded from a dedicated property tax, with a $2.8 million maximum levy. The economic model for this plan includes assessed value, real market value and the impact of constitutional property tax limitations on the available levy.

- *Consistent*. Although Ralph Waldo Emerson said "a foolish consistency is the hobgoblin of little minds," it is appropriate for an economic model to avoid gross inconsistencies. For example, the inflation model underlying interest rates, bond returns and property tax levies should be the same in each case. This means that pieces of different models should be combined only with great care. Consistency can be overdone. While, for example, it may be generally true that dividend yields are correlated with inflation, to force such a relationship to occur in an economic model is another matter entirely. Over-specified models may be too narrow, without enough variability to cover all possible future scenarios.

- *Stable*. If the economic model is used repeatedly, the results it produces generally should not vary a great deal from one application to the next. For example, if a model produces a real return on equities of 6 percent in one study and 3 percent in a study a year later, there should be some compelling explanation for the changed results.

- *Appropriate*. The economic model should be appropriate to the task at hand. Most of our models are focused on the long term—30 years or more—and these will be the ones used in this discussion. However, it is easy to imagine that a plan sponsor could want to know the probability that the cost of the plan could exceed a given threshold within the next five years. In the latter case, a completely different economic model may be appropriate.

- *Simple*. It's important not to go too far. An overly complicated economic model becomes an impediment rather than a tool. At least two issues are involved: Communication and expense. Obviously, simple models are easier to explain to our clients. Most critically, a simple model makes clear the limitations of all modeling: Clients are made aware that the economic model is far less complicated than the economy itself and that all projections are highly uncertain. It is far wiser to make the uncertainty explicit to our clients than to obscure it within technical jargon and complexities. Furthermore, a simple model can draw clients into the modeling discussion. They can understand the construction of the model and participate in setting parameters. We want cheap economic models that we can use on all projects. Committee-designed economic models, constructed in the hope of having one model to fit all plans, can become impractically complicated. Such models can become expensive to run, which impedes their use in any but the best-funded projects. Furthermore, adding variables to such a model—for example, property tax revenue as discussed previously—can be prohibitively expensive.

- *Dumb*. In building an economic model, it is wise to avoid prediction. Instead, focus on what could happen, rather than on what will happen. For this reason, I favor models that are purely mechanical in construction, and that avoid any knowledgeable economic input on the part of the actuary, investment managers, sponsor or board. In my experience, models that result from Delphi processes among knowledgeable professionals tend to be both optimistic and narrow, producing higher and less variable returns than is ultimately the case. Moreover, a purely mechanical model is testable.

- *Testable*. It would be wise to test our economic model, to ask what it would have projected a number of years ago and where the actual economic results would have fit in the range of values projected by the model. Below I present and test an algorithmic economic model we have used in a number of studies.

Traditionally actuaries have used a simple deterministic asset model. Single rates of inflation, salary escalation and asset returns were selected, based on long-term trends and projections. This economic model isn't all bad: It can be comprehensive enough in most cases, and it is certainly consistent, stable and simple. Moreover, it facilitates the analytical approaches to plan funding that are commonly used in the profession.

However, a deterministic economic model has at least two major drawbacks:

- First, it does not convey to the plan sponsor the *range* of possible plan costs that may arise in the future.

- It facilitates the use of discounted present values, often labeled as liabilities, and presents these as point estimates of the funding status of the pension plan. This has resulted in ill-advised decisions on the part of plan sponsors and boards.

Developing a stochastic economic model is not that difficult as long as the purposes of the model are kept in mind. An airplane must fly in many kinds of weather, but the Wright brothers did not need to predict the weather to build an airplane. The point is that our economic model need not predict the economy, it only needs to subject our plan to a wide range of reasonable economic scenarios so we can measure the plan's performance in each, and in the aggregate.

The economic model we will use in this example is a simple stochastic model based on a statistical resampling technique known as the "bootstrap" (see Efron and Tibshirani 1993). We proceed as follows:

- We select a base period for resampling. Typically, I use the last 30 years of economic data, including rates of inflation and total return on major asset classes.

The base period selection is arbitrary in some ways: We want a recent period of sufficient length to offer a wide range of economic scenarios, but short enough to be reasonably current.

- For each simulation trial, we select a series of random five-consecutive-year periods within the base period. The five-year periods are selected with replacement, so the same or overlapping periods may be selected multiple times in each trial. We wrap the base period around, last year to first year, so that a five-year period selected within five years of the end of the base period will get its remaining years from the beginning of the base period.

- The five-year periods are successively concatenated into an economic scenario for the simulation trial. Annual inflation rates and returns are taken from each year in the five-year periods.

- The resampling process is repeated for each trial.

This economic model meets a number of our criteria. It is simple, mechanical (dumb) and testable; it does not depend on knowledgeable projection of future economic trends nor on sophisticated statistical analysis of past data. The model can be made comprehensive: If adequate historical data is available, the approach can be extended to include the yield curve, property values, income tax receipts and other variables of interest.

Consistency is enforced by the use of actual past results. The resampling model is generally appropriate for long-term projections that are at least several times the resampled block of five years length.

Interdependencies of asset class returns and inflation—both serially and within individual years—are preserved (mostly) by taking blocks of individual yearly returns.

This model can be back-tested. Here's one such test:

- Select an asset mix. In this case we use a typical mix of 60 percent equities and 40 percent fixed income.

- For each 30-year base period in the database (starting in 1926 here), simulate the average return for the 10 years following the base period.

- Determine where in that distribution of average 10-year returns the actual returns for the next 10 years falls.

For example, for the 30-year period from 1950 through 1979 we would simulate the average return for 1980 through 1989 and then compare the actual average return for that 10-year period with the simulated distribution.

Figure 1 plots the percentile of the actual 10-year return within the distribution of simulated returns for 30-year periods ending in 1955 through 1992. In only two cases does the actual return lie outside the distribution: In 1980–1989 and 1982–1991. Therefore, this model at least produces enough variability to contain actual results.

Figure 1

Actual 10-Year Average Returns vs. Simulated Returns

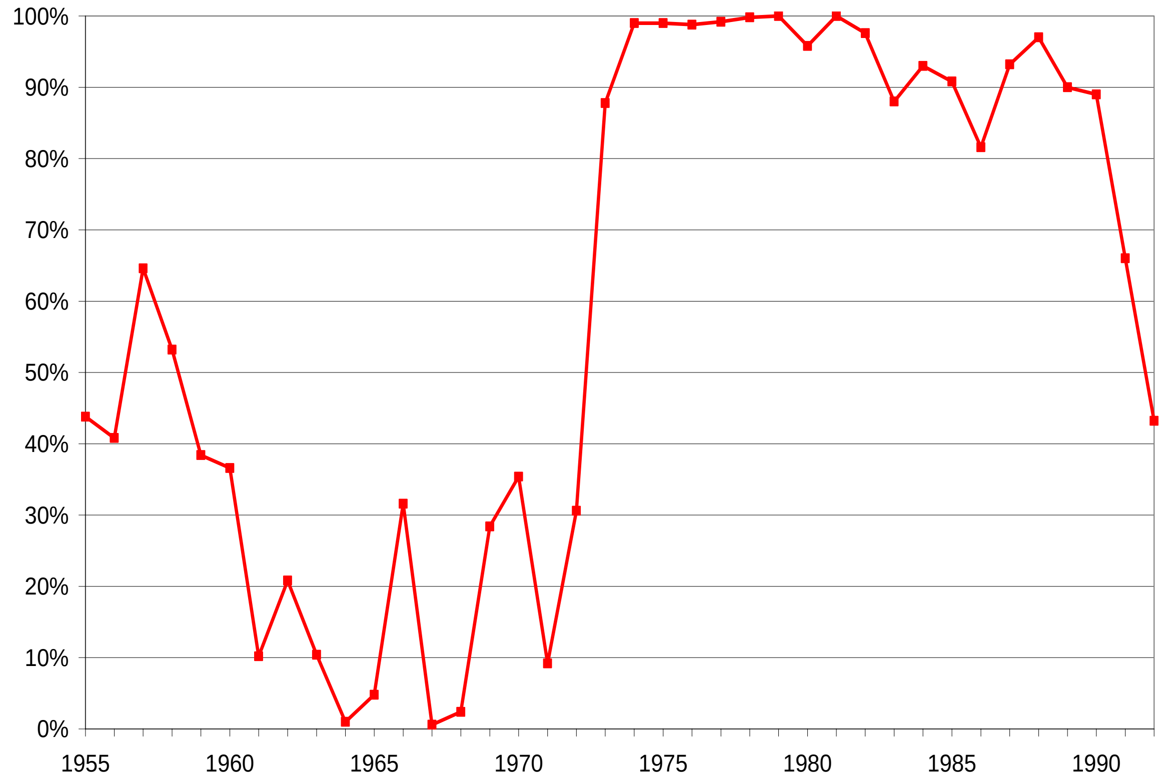

Note: The percentile rank of 10-year average returns is shown within 500 simulated 10-year returns based on resampling from 30-year periods ending with the year shown.

This is only one test. I invite you to invent your own.

One weakness of the resampling model is stability: When the resampling period of 30 years is updated one year, one year of experience is added and one year is removed from the data available for resampling. If, for example, a year of high returns is added and a year of low returns is removed, the results of the economic model can be materially affected.

Nothing is perfect. This is only one possible model. I invite you to invent your own.

Armed now with a simple stochastic economic model, we can proceed with our benefit projection.

2.4 Project Plan Benefits (for Real This Time)

We are now in a position to proceed with our projection of future plan benefits. Clearly, our plan model and our economic model must be integrated: As our projections proceed into the future, inflation driven increases in salaries and benefits must depend on the inflation results produced by the economic model.

There may be other ways in which the economic model may affect the projection of benefits from the plan model. For example, it is not uncommon to find benefits that depend on "excess" investment returns. COLAs that depend on an accumulation of investment returns over the actuarial assumption are particularly common.

Based on the plan and economic models described previously, I have projected the benefits payable from the example plan. Results are shown in Figure 2.

Figure 2

Projected Plan Benefits as a Percentage of Active Member Payroll

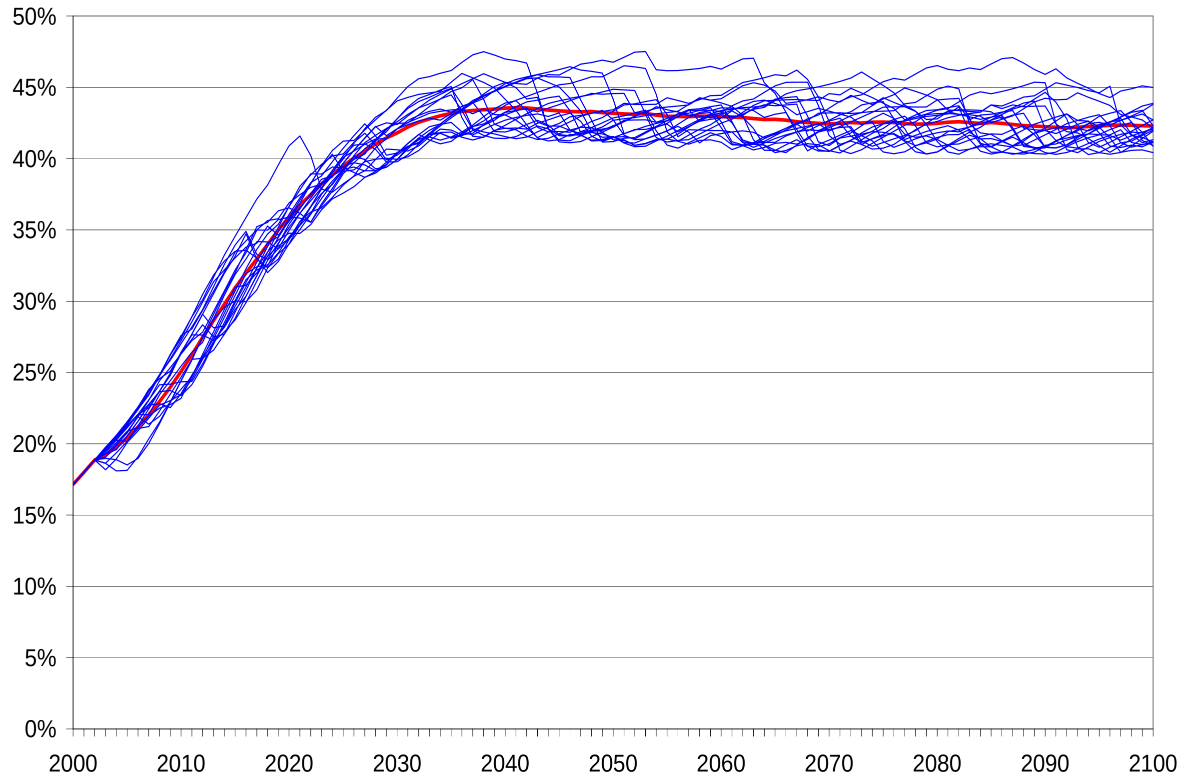

Note: The results of 14 of 100 individual simulation trials are shown in blue; the heavy red line is the average for each year.

What we see in Figure 2 is typical of many plans we have studied:

- The plan is fairly immature: Benefits are increasing in relation to active payroll. In this particular case, benefits have recently been improved.
- Benefit payments often show a modest peak around 2020 to 2040 or so. This is the impact of the postwar baby boom.

2.5 Decide on a Funding Strategy

Given that we now have a projection of plan benefits, we can start thinking about how we want to pay for them. (Note here that we will discuss real contributions rather than accounting expense.) We have a wide range of choices:

- *Traditional actuarial funding*. Here we attempt to fund benefits as a level percentage of pay over each member's working lifetime. There may be an initial period during which we have to make extra contributions to make up for past contribution shortfalls, benefit improvements or asset losses.

- *Funding relative to a revenue stream*. We may want our plan contributions to be a level percentage of some revenue stream, such as property tax collections, as discussed previously. This revenue stream need not be smooth. For example, transit districts frequently have temporary local, state and federal subsidies—both for capital outlays and operations—that expire periodically. Such temporary revenue sources may suggest corresponding increases in plan funding while they exist.

- *Funding to a target cost*. We may want to accumulate assets over a given period so that the ongoing contribution after that period is some desired percentage of pay.

- *Funding to a target proportion of cost*. One important reason for funding a pension plan is to have the investment markets pay a portion of the ongoing plan cost through returns on assets. For example, we may decide to accumulate over 30 years a sufficient asset pool so that expected investment returns equal 75 percent of the plan outflow, measured in terms of benefits and expenses.

All of the previous approaches involve setting an asset target and funding to reach it. Note that "liabilities" have not been discussed. For purposes of this example, we will opt for traditional actuarial funding.

2.6 Project New Hire Cost

Given that we have decided upon traditional actuarial funding, we want the cost of the plan at some point to be the cost to fund a new member's benefits over his or her working lifetime. So, as a first step, we determine the cost of a new hire.

We proceed as follows:

- Determine an average group of new entrants. A typical plan model will involve a number of more or less homogenous populations; generally, the criterion for a population is that the members share the same demographic characteristics and assumptions. In studying new entrants, I typically look at the new members joining each population over the past few years and derive somewhat smoothed distributions of age and entry pay.

- Project future benefit payments to these new entrants. Note that this projection may occur once—if the economic model is deterministic—or once for each simulation trial in a stochastic model.

- Simulate the active payroll of these new entrants. Again, the payroll projection may be done once for each simulation trial.

- Compute the new entrant cost. For each simulation trial we determine the level contribution of payroll necessary to pay benefits for the aggregate new entrant population in such a way that all assets are exhausted by the last payment to the last new entrant alive. Typically, a 100-year projection is used. Note that a present value calculation is not necessary in this determination. In fact, given a stochastic rate of return and inflation in each trial, a present value calculation could be an unnecessary distraction. Instead, the computer determines the assets remaining at the end of the projection—usually negative—with a 1-percent new entrant contribution and with a 0-percent new entrant contribution. Since the final asset balance is a linear function of the new entrant contribution, a simple ratio then provides the new entrant cost. The computer can do this calculation quite rapidly.

The result of our efforts is the distribution of new entrant costs shown in Figure 3.

In passing, note that the new hire cost computed here may be a useful figure for the plan sponsor. The computation of the new hire cost described here does not use the usual entry age normal approach: Instead, it depends on recent hiring patterns, rather than those that prevailed in the past. Furthermore, it is only the current tier of benefits that is included in the new hire cost computed here.

Figure 3

Simulated New Entrant Costs for the Example Plan

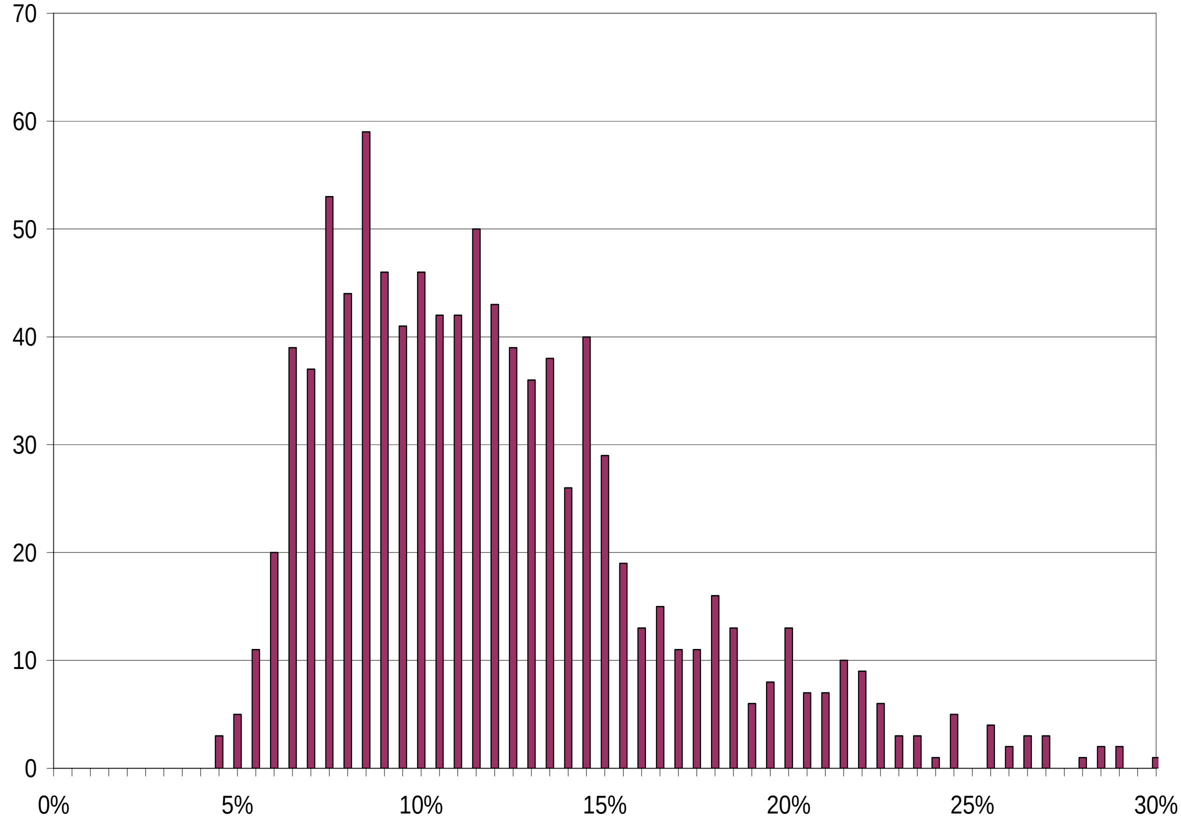

Note: The distribution of total employer and employee cost as a percentage of payroll for 1,000 simulation trials is shown.

Figure 3 ought to scare the pants off every pension actuary for two reasons:

- *Variability*. We would typically inform a plan sponsor that the new entrant (normal) cost of their plan is around 10 percent of pay for this plan. In actuality, it could reasonably be between 5 percent and 30 percent of payroll.

- *Skewness*. Note that the graph is skewed positive. Based on a symmetrical distribution of asset returns, the distribution of the cost of a deferred annuity will always be skewed in this way because the graph of annuity cost versus rate of return is concave up. This means that when we set our actuarial assumption based on the mean of a more or less symmetrical distribution of expected asset returns, the cost we provide our clients is the mode or median cost. Because the cost distribution is skewed, the mean cost is actually higher than the one we typically calculate. To provide our clients with the mean cost, we would need to reduce our assumed rate of return by anywhere from 0.25 percent to 1 percent, depending on the variability of plan assets.

But I digress.

We now have a distribution of new hire costs, and in particular, a new hire cost for each simulation trial.

2.6.1 Select an Amortization Period

Here we select a number of years after which we want the plan cost to be the new entrant cost. For purposes of this example, all future contributions, during and after the amortization period, will be computed as a level percentage of active payroll. In this example, we will use 20 years.

2.6.2 Compute an Ultimate Asset Target

At this point we need to set an end point. We are going to assume that after the end of our projection—100 or more years in the future—the equation of balance holds: Specifically, contributions plus the real return on plan assets equals benefits plus expenditures.

We will assume that after 100 years, benefits and other expenditures have stabilized at a constant percentage of active payroll. This percentage could be the average over the last few years of the projection or it could be simply the outflow in the last year of the projection. We set our desired plan contributions to be the new entrant cost computed previously.

It remains to estimate the return on plan assets after 100 years. We use the average real return on plan assets during the 100-year projection period. Based on this assumed return, we can compute the level of assets needed at the end of the projection period for the equation of balance to hold thereafter.

Note the asset target is computed separately for each simulation trial.

2.6.3 Compute an Asset Target at the End of the Amortization Period

Given that:

- The plan contribution rate from the end of the amortization period to the end of the projection period is to be the new entrant cost,
- Benefits and expenses are to be paid from the end of the amortization period to the end of the projection period,
- Earnings on plan assets are as simulated by the economic model, and
- A target level of assets is to be achieved at the end of the projections period,

we can compute a unique asset level that needs to be attained at the end of the amortization period to make all of the above conditions hold. This becomes our funding target at the end of the amortization period; we usually express this target as a percentage of active payroll.

Note again that this target will be different for each simulation trial.

2.6.4 Compute an Initial Cost for Each Trial

Given that:

- A target level of assets is to be achieved at the end of the amortization period,
- Benefits and expenses are to be paid during the amortization period,
- A current level of assets is on hand, and
- Earnings on plan assets are as simulated by the economic model,

we can compute a unique contribution level that needs to be maintained from now until end of the amortization period to make all of the above conditions hold. This becomes our current contribution level; we usually express this target as a percentage of active payroll.

Note that this contribution level will be different for each simulation trial.

The result of our efforts is shown in Figure 4. The average contribution is –3.49 percent for the first 20 years and 6.31 percent thereafter. We note that the average employer contribution for the first 20 years—our amortization period—is negative because the plan was overfunded. That was the situation as of July 1, 2000; regrettably, it is no longer the case.

In this example, we have allowed the negative contribution to show the impact of overfunding. In most actual cases, we would have constrained costs to be nonnegative.

3. What Have We Done?

It's worthwhile to stop and take stock. What have we done so far?

- We have created a simulation model of our example pension plan. Our model consists of a plan model for generating cash flows and an economic model for simulating future economic scenarios in which those cash flows will be generated.

  The combination of the two models produces an open-group, stochastic simulation model of the retirement plan.

- From first principles, not using calculated present values or liabilities, we have designed a reasonable way of funding this retirement plan. Our funding method proceeds by analyzing and matching cash flows: A level percentage of payroll from the employer plus simulated asset returns minus simulated benefits and expenses. This type of cash-flow matching could be applied to employer contributions based on property or income tax levels; it can easily be generalized.

Furthermore, our model naturally determines a distribution of the cost of the retirement plan. The plan sponsor and the actuary are presented with a range of probable costs, not merely a point estimate.

Figure 4

Simulated Total Employer Costs for the Example Plan

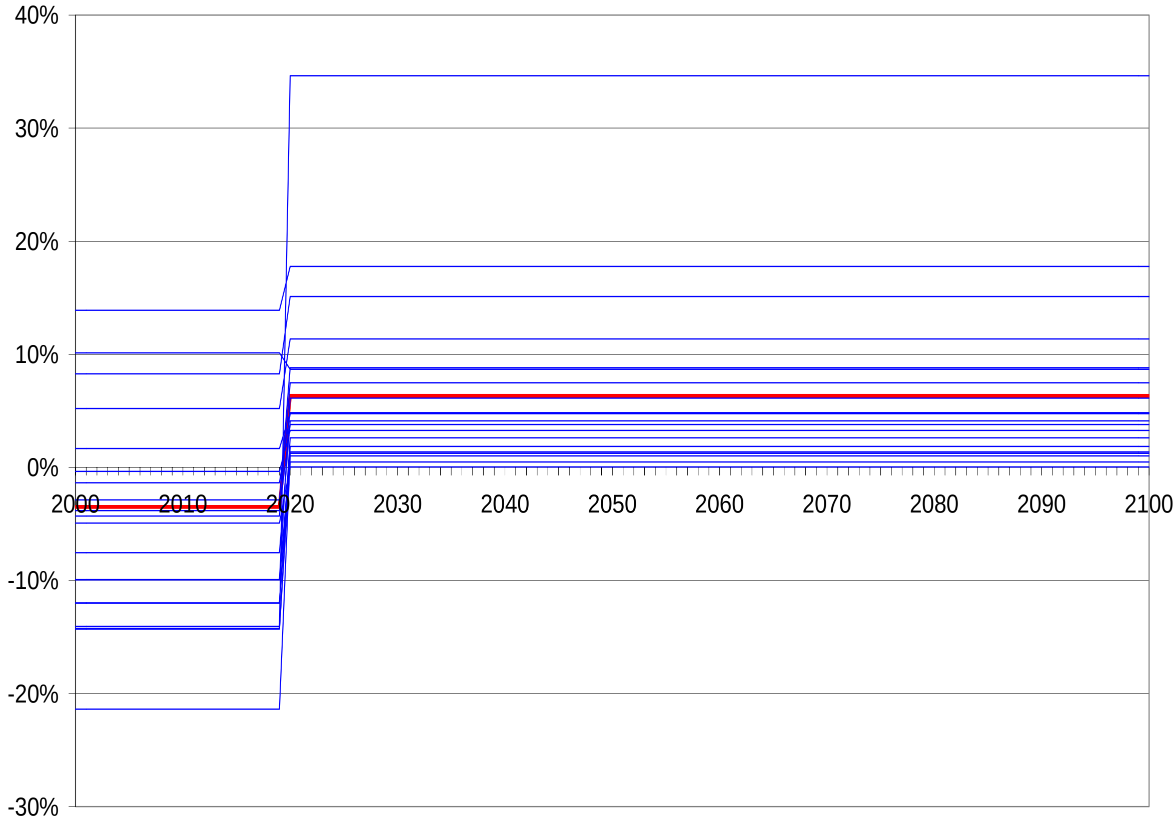

Note: The results of 14 of 100 individual simulation trials are shown in blue; the heavy red line is the average for each year.

It's also worth noting what we have *not* done. Note that we never selected an assumed rate of return on plan assets. While some may assert that building a credible economic model is more difficult, I feel that nonparametric approaches, such as those described here, can produce credible and testable economic models that are more than adequate for our purposes. Moreover, any stochastic economic model, no matter how simple or even flawed, is better than an assumption that plan assets will earn a fixed, unvarying rate of return forever.

4. What are the Advantages of This Approach?

4.1 Complete Models Are Constructed

We leave a lot out of our traditional actuarial models. Variation, for a start. Including variation can be particularly important when some plan benefits depend on economic variability for their very existence. Investment-driven cost of living arrangements are an example.

In addition, traditional closed-group models omit some important population dynamics, such as the interplay of various benefit tiers that is often found in public plans. Furthermore, we may be able to reliably project near term growth (or decline) in the active workforce due to annexations, extensions of service or layoffs. Such events can have a material—and largely predictable—impact on the cost and funding status of a plan.

The open-group, cash-flow matching approach described here allows us to model all relevant aspects of a pension plan, including future events that are material and can be predicted with a fair degree of confidence.

4.2 No Liabilities Are Computed

Proceeding from first principles, we have not created an actuarial accrued liability.

The "liabilities" of a pension plan have been attracting a great deal of ill-informed commentary among accountants, financial analysts and investment advisors. In particular, many have opined that the "liabilities" of a pension plan—often defined no further—should be valued at some kind of "market value".

The prior simulation shows a pension plan for what it is: an ongoing, dynamic financial system. In the example, we have funded this system by matching the incoming and outgoing cash-flow streams.

It is important to remember that liabilities do not exist in the physical world. Instead, they are abstractions, mental tools that allow us to organize and deal with the realities around us. To the extent they mislead or confuse, liabilities should be set aside, just like any other tool that is not suitable for its job.

4.3 We Facilitate a Clear Definition of Any "Liabilities" That Are Computed

If the actuary is asked for the "liabilities" of the plan, the question to ask back is the purpose for which the liability figure is requested. If the purpose of the liability is to measure the market value of the plan sponsor's obligation, then, in my opinion, the only suitable liability is the plan's termination liability.

The only circumstance in which the plan would be forced to liquidate all of its assets and liabilities at market value would be in the event of a plan termination.

In the private sector, the plan's termination liability is well-defined and computable. To avoid a full computation of PBGC termination liability, an approximation could be the vested accrued liability using PBGC or similar rates. Any excise taxes on the reversion of assets to the plan sponsor should probably be included.

In the public sector, the situation is far less clear. Depending on the state, a plan sponsor may be able to terminate a plan, or it may be required that any termination or reduction in benefits apply only to future hires. In the latter case, it is unlikely that the plan sponsor would discharge its obligations by liquidating assets and buying annuities at market value. It is far more likely that plan funding would continue as before. Given that the last obligations of the capped plan could be discharged more than 50 years in the future, it is unclear that market conditions today have any relevance at all.

4.4 We Distinguish Between Liabilities and Funding Targets

The actuarial accrued liability is not a liability at all: It is a funding target. The actuarial accrued liability is the assets that would have been on hand had we been contributing the normal cost all along and actuarial assumptions had been exactly met. To the extent that our assets are below this target, we have to contribute an extra amount to catch up. If our assets are in excess of the target, we reduce our contribution rate.

A funding target can be computed in the above example. For each simulation trial, we can compute the amount that would need to be on hand today so that the plan cost during the amortization period exactly equals the new hire cost. A distribution of the results of such a calculation is shown in Figure 5.

Figure 5

Distribution of the Funding Target for the Example Plan

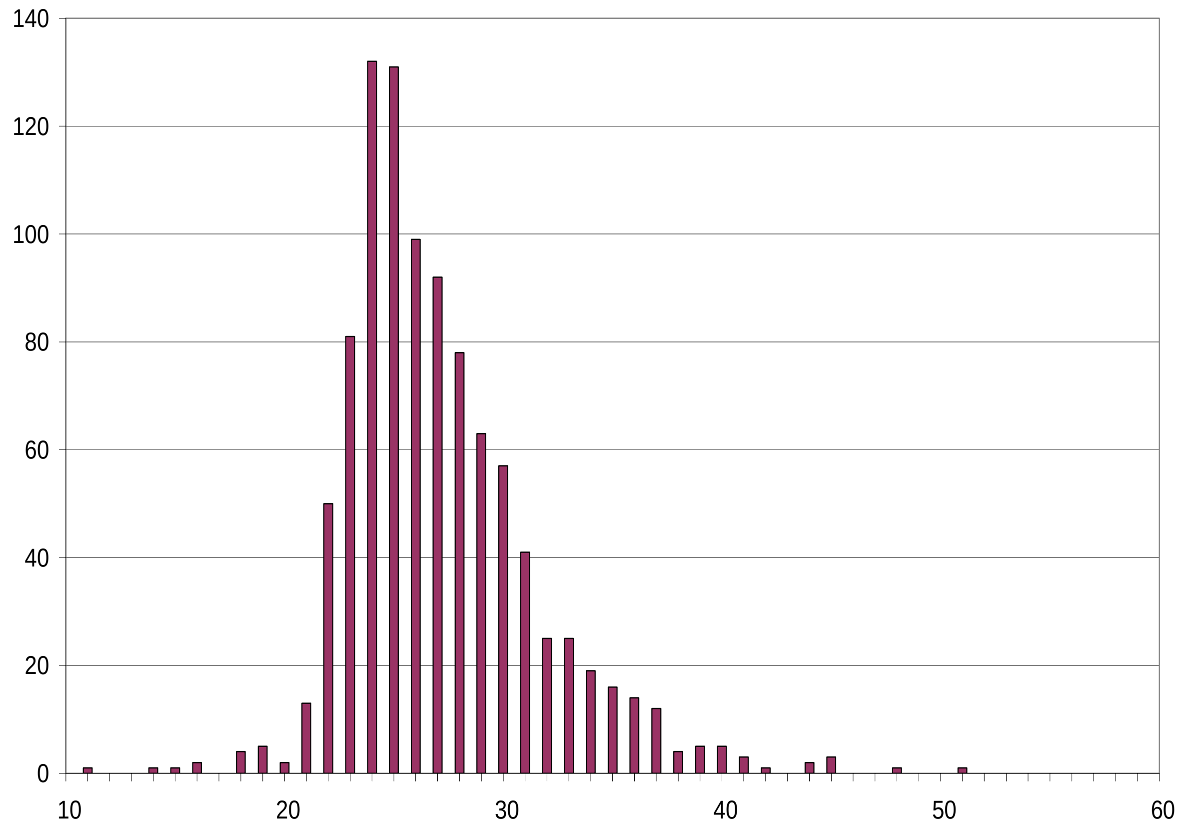

Notes: The horizontal axis is assets at time 0 in billions of dollars. The distribution of 1,000 individual simulation trials is shown.

Note that the funding target is expressed as an asset level desired to achieve a certain result. There is no reference to present values or to liabilities.

4.5 Funding Ratios Are Presented as a Distribution

The ubiquitous funding ratio has caused a great deal of harm. I have seen plan sponsors and boards decide to increase benefits substantially based on the "excess" assets in the pension plan. Unfortunately, the investment markets had their own ideas about what to do with the excess assets.

Any point estimate of a probabilistic entity is bound to cause trouble. A far better approach is to present plan sponsors with an appropriate distribution. Any distribution, based on any stochastic economic model, is better than a point estimate.

Figure 6 shows the funding ratio distribution for the example plan.

Figure 6

Distribution of Funding Ratios for the Example Plan

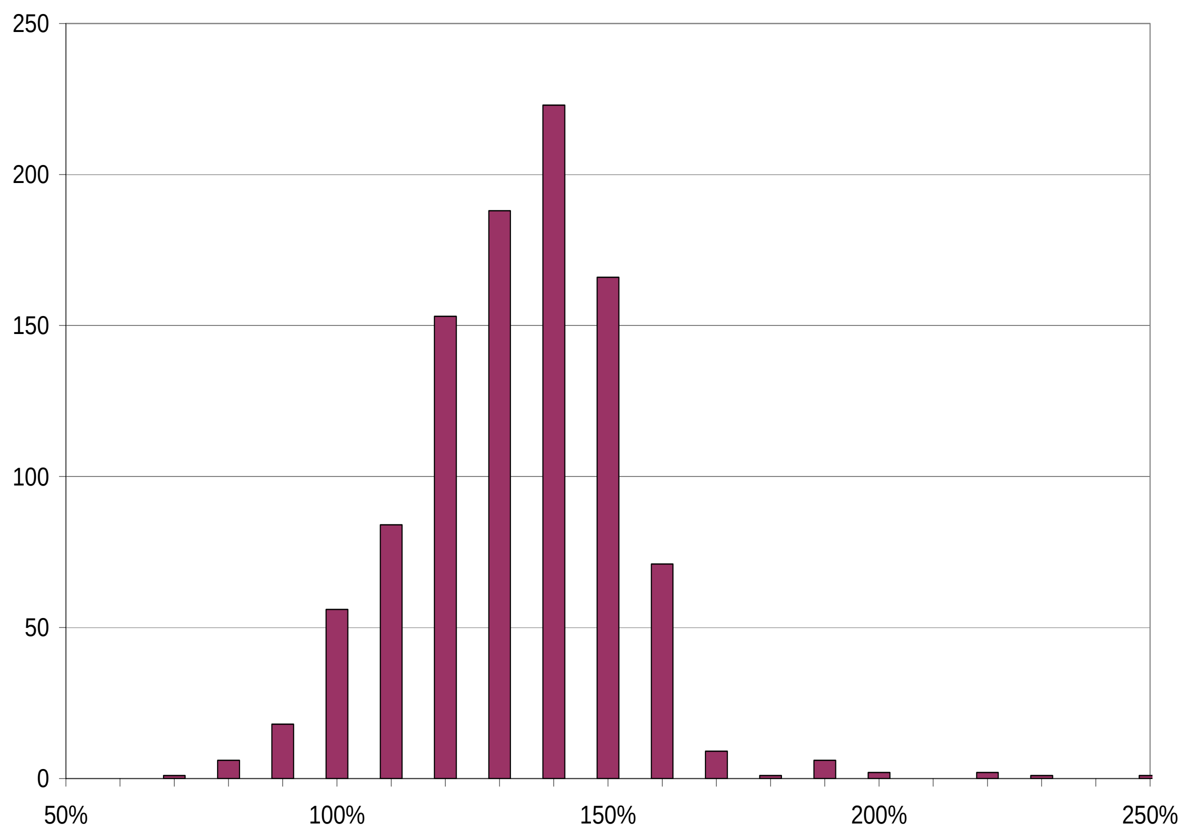

Note: The distribution of 1,000 individual simulation trials is shown.

Note in Figure 6 the enormous range of the funding ratio. While the average is 121 percent, it can reasonably be said to range from 80 percent to 170 percent. Decisions based on this information will be based on better knowledge of the uncertainties involved in pension funding.

4.6 All Results Are Distributions

As mentioned previously, presenting a result as a distribution is both more informative and less likely to cause confusion than offering a single point estimate—even if it is a best estimate. Under the approach presented here, all results—costs and funding ratios—are presented as a mean with an associated distribution.

4.7 Projection of Results Is Natural

Since the model is designed to be an open-group, stochastic model, projection of the results of future studies is easy and natural.

For example, it is a straightforward addition to the model to incorporate the results of an annual closed-group, deterministic, liability-based actuarial valuation in each future year. In this way, the traditional actuarial cost can be projected and analyzed stochastically.

Figure 7 shows the simulated closed-group actuarial cost for the example plan.

Figure 7

Simulation of Traditional Actuarial Valuations for the Example Plan

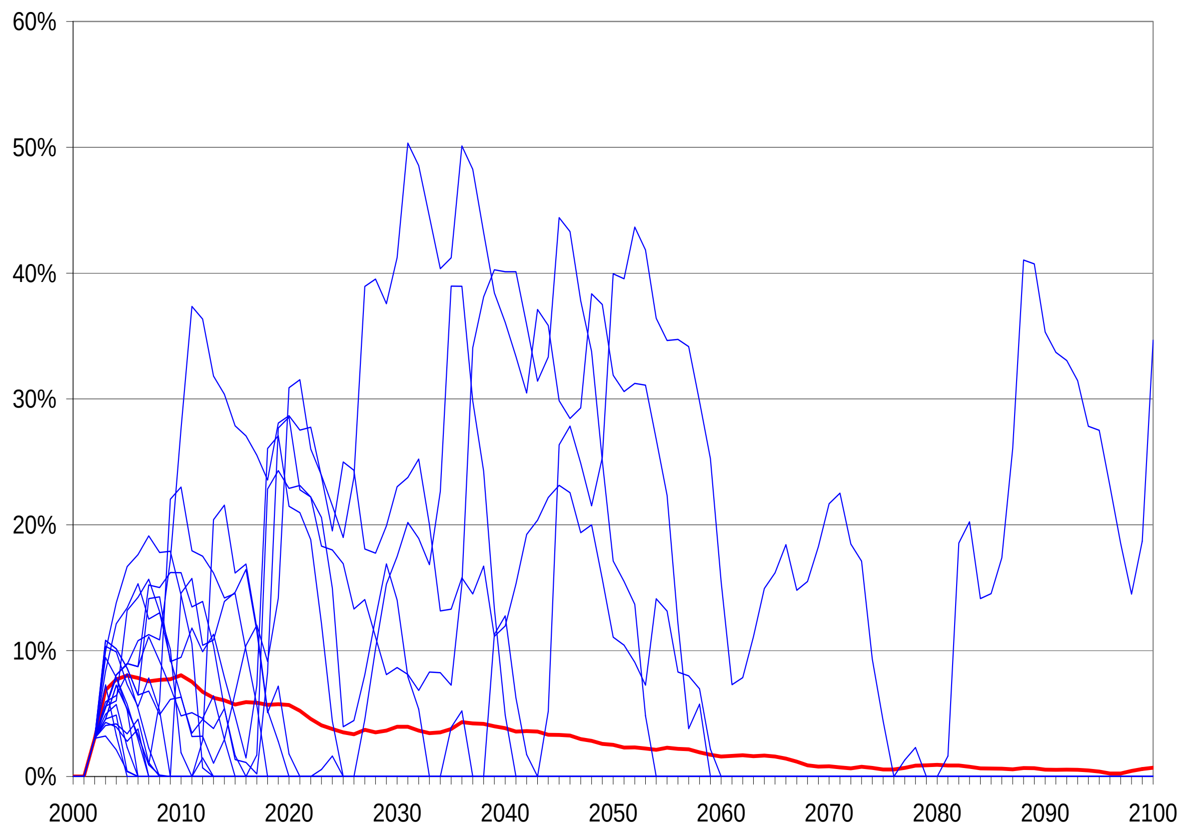

Note: The mean and 14 of 100 individual simulation trials are shown.

We note in Figure 7 that our traditional methods—in this case, entry age normal—don't do such a bad job of producing reasonable and reasonably level plan costs, at least on average. Whatever their faults, current funding methods were invented by smart and capable actuaries, who deserve our appreciation.

4.8 Our Results Are Understandable

It has been my experience that plan sponsors and boards understand simulation and distributions of results far more easily than even the basics of entry age normal funding. Simulation, in particular, is all around us; there are probably relatively few of the people with whom we deal that don't have a flight, driving or sports simulator lurking on one of their home or work computers.

Once the graphs are explained, the presentation moves to substantive issues quite quickly. Furthermore, the transparency of simulation methods allows our clients to participate fully in all discussions and to absorb the information they need to make the decisions for which they are responsible.

4.9 Experimentation With Pension Dynamics

An efficient model allows the actuary to experiment with the dynamics of the pension plan, easily answering "what if" questions. For example, how does the funding status of the plan affect the variability of results? What if inflation decreases further? How does the lower tier-2 benefit level affect future costs?

This sense of experimentation and play is most important to the actuary as we analyze our results and formulate our presentation to our clients.

4.10 Easy Segue to Asset Allocation

Investigating the impact of asset allocation is easy and natural using an open-group stochastic model, especially when closed-group, deterministic valuation results are included.

Figure 8

Simulation of Traditional Actuarial Valuations for the Example Plan

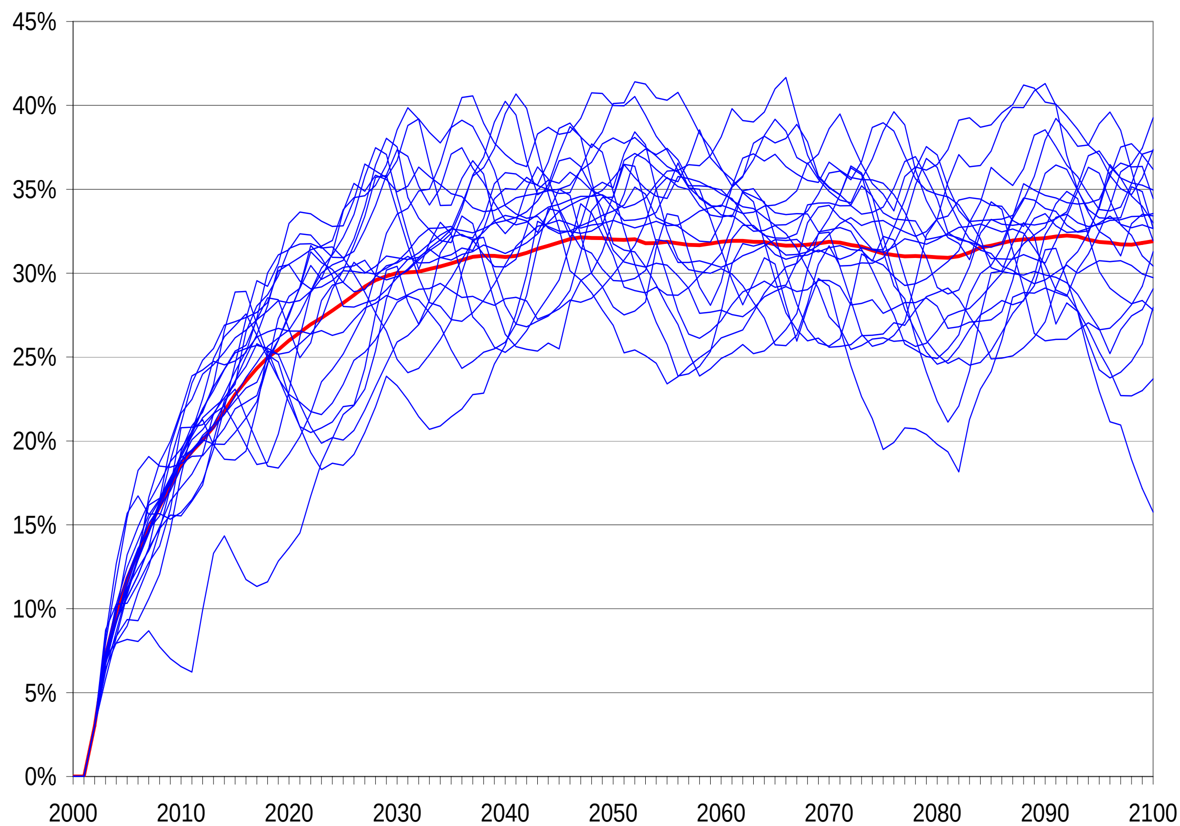

Notes: Assets are invested entirely in cash equivalents. The mean and 14 of 100 individual simulation trials are shown.

For example, Figure 8 shows the projected future closed-group plan costs of the example plan if all assets were invested in cash equivalents. In this example, the plan actuary doggedly keeps his economic assumptions unchanged—waiting, no doubt, for common sense to return to the investment committee—so actuarial losses accumulate, gradually driving costs up.

5. What are the Disadvantages of This Approach?

The approach I have outlined does not come without costs. Here are a few.

5.1 Regulatory Restrictions in the Private Sector

In the private sector, the slow-motion regulatory destruction of DB pension plans continues apace. In particular, the practice of closed-group, seriatim, deterministic actuarial valuations is ossified in the Internal Revenue Code and accompanying regulations. Accordingly, open-group, cash-flow matching approaches cannot be used.

There is still room in the private sector for simulation studies to augment the actuarial valuation, and many firms are working in this area. However, the funding complexities introduced by ERISA do cause significant complications—and computational penalties—for private sector models.

5.2 Construction of an Economic Model

Many actuaries may feel hesitant about constructing an economic model. They shouldn't be. I have discussed one possible model. An infinite number of other models can be constructed.

The key point is to experiment with multiple models. Build two or three radically different models—some parametric, some nonparametric—and experiment with the plan. If the results look similar with several different models, then you may have robust results. If the projections change radically from model to model, you have found something worth investigating: What is changing? Why? What does it mean? What does it mean in terms of decisions you and the plan sponsor must make?

An economic model need not be—in fact, must not be—the kind of complex, predictive model used in the investment field. Our work has different needs.

5.3 Budgetary Limitations

The modeling described here takes time. It may be too time consuming—and expensive—for medium and small plans.

However, much can be accomplished with a little ingenuity. Our computers are now fast enough that even simple tools can do great things. Recently I wrote an Excel spreadsheet that performed an entry age normal valuation of a government plan with about 20,000 active and inactive members. The recalculation time was about eight minutes. I've run stochastic valuations of small plans with spreadsheets as well.

The point is that interesting things can be done with simple tools and fast computers.

5.4 Conflicts With Other Professions

The accounting, financial analysis and investment professions have grown accustomed to our liability calculations and have based much of their professional activities on the results we produce. Changes on our part will not be welcomed.

However, it is important to recognize that many of the accounting rules currently in force, especially in the private sector, may misrepresent the true impact and risk of DB plans on sponsoring companies. This is not an area in which I am an expert, but it is vitally important.

5.5 New Systems

The approach outlined here will require the development of new software. The software required is neither simple nor easily developed. Naturally, programming expertise is required, but even more important is actuarial expertise to identify which aspects of a pension system must be modeled and which can be omitted.

In the simulation community, the phrase is "precision for decision": Make the model as complicated as necessary to address the decisions that need to be made, but no more complicated.

6. An Empirical Approach

Actuaries are trained mathematically, so we approach problems analytically. We use equations to produce best estimates.

The problem with this approach is that pension plans are neither mathematical nor physical entities. No physical laws govern their behavior. No differential equations apply. Accordingly, analytical techniques may be inadequate to the task.

I suggest a more empirical approach to retirement systems. Regard them as an interesting creature of unknown behavior, build a computer model of the system, study the predicted behavior of the system and, when possible, compare the predicted behavior with actual behavior.

If we go straight to modeling, the important and analytically difficult features of the plan—new entrants, benefit tiers, investment-based benefits, economic uncertainty—will be fully recognized in our work and in our advice to our clients.

7. Our Own Voice

Above all, the actuarial profession needs to find its own voice. We are told what to calculate and even how to calculate it by Congress and state legislatures, by federal and state regulators and by the accounting profession. The requirements imposed on us are complicated, contradictory and, in some cases, just plain crazy. It is time for our profession to determine—based on our own expertise and standards—the best way to model pension plans and to insist that regulators and other professionals accept our standards.

Reference

Efron, B., and R. Tibshirani. 1993. *An introduction to the bootstrap*. New York: Chapman and Hall.