What it Really Means to Use a High Discount Rate for Pension Liabilities

Avoiding the next pension funding crisis

M. Barton Waring[1]

Draft of Feb 17th, 2012

Almost everyone is aware of the scope and depth of the pension funding crisis in the US, with public employee and Taft-Hartley pensions hovering in the range of a 65 or 70 percent funding ratio, and with corporate plans around 85 or 90 percent funded. In all cases contribution rates are pushing to the limits of the tolerable—and in the case of many plans even beyond. Precision is not needed; just focus on the plan most important to you—you know its numbers, and they aren't reassuring.

But not everyone is familiar with how we got to this sad place, and the role that the actuarial method for setting the discount rate—using the expected rate of return on the assets—has had in getting us there. And if we're going to avoid future pension funding crises, then board members, pension staff, sponsor executives, employee representatives, and the actuaries and accountants involved (and their regulators) need to understand just what is wrong with the current discount rate methods.

First, where do contributions come from?

When funding ratios are discussed, they usually are *accrued liability* (AL) funding ratios. Let's remember that the AL is just what it says it is, the *accrued* portion of the larger "present value of all foreseeable benefits for current and past employees"—the anticipated total deferred compensation-pension debt to the employees. This larger and more important value is usually referred to in short form as the *present value of future benefit payments* and sometimes just as the present value of benefits (and we'll abbreviate it here as the PVB). Each period a *notional* "payment" is made on this larger PVB "debt", and this payment accrues to the balance sheet as an "accrual accounting" addition to the accrued liability (and it simultaneously accrues to the income statement as a component of pension expense). This AL represents the portion of future benefit payments that has been "earned" up to this time by the employees, at least in some rough sense. So you can see that the normal cost, and thus also the contribution is a function primarily of the PVB, and secondarily of the AL.

I used the word "notional" to describe normal cost because, for purposes of building up the AL, it is irrelevant whether or not it is accompanied by a cash contribution—it is used solely to build up the

[1] This paper is copyrighted by the author, and is not yet final and thus subject to ongoing revision. It has been submitted to the Financial Analysts Journal for possible publication there.

tracking version of the liability known as the AL, and as a component of pension expense—*it does not yet represent any exchange of cash.*

The notional normal cost payment is exactly like a scheduled mortgage payment or, even better, like a regularly scheduled payment to build up a college fund for one's child, with a specific target value in mind for a specific future date, when the child will graduate from high school. You can look up the payment needed for each individual employee on an amortization table, or compute the amount needed on your HP-12C® financial calculator (on a level payment basis). Added up across all employees, you get the aggregate normal cost.

And of course cash (*actual* cash, not notional!) contributions should be made in that same amount, normal cost. If that is done each period, and if the assets being built up earn a return equal to the discount rate, the AL is always fully funded. Of course, if some loss happens to the fund built up from these contributions (such as an investment loss), then the AL is no longer fully funded.

The point is that the AL is just a portion of the full underlying liability, the PVB; that portion represents a point of progress in the payment amortization of that debt over the employees' working lives, a sort of funding target protecting the security of earned benefits. A 100% AL funding ratio means that the portion already earned is fully funded, so that earned benefits are secure. But by the same token it also means that the full liability, the larger PVB, is *underfunded*. This is the intent, after all—we're not providing full funding initially, promising instead to fund the plan only as benefits are earned and accrued.

And it then must also be true that if we're underfunded on the AL, we're far more underfunded than we should be, and it is a serious matter. We're way behind on our mortgage payments, and interest on the full PVB is still being added to the principal balance. An additional contribution, over and above normal cost, should be made to make up for this shortfall.

But just as it is hard for folks to catch up when they get behind on their mortgage, the further behind you get in funding the pension plan, the more difficult it is to figure out how to catch up. Payments in the future will have to be larger, often much larger, than if we had made them on time. It is not a good thing to be behind on the mortgage payments, and it is equally not a good thing to be underfunded relative to the AL in a pension plan. It isn't a casual matter, although it often seems to be treated that way, "we're in pretty good shape, we're 85% funded."

But a good question for all of us to ask ourselves is, "how did it become acceptable and even tolerable to be underfunded on the accrued portion of the liability?"

Contribution Calculations, and some self-help for the concerned non-actuary

It will be helpful to our discussion if we work through just what it means to calculate and make a contribution. This analysis will also have a side benefit, of suggesting some easy ways for lay people with a very modest amount of finance knowledge to calculate the contributions that will be needed by a given plan. This is especially useful when a plan is underfunded, and it can serve as a good check on the actuarial advice the plan is receiving with respect to its contributions.

One approximate approach to an aggregate version of the *normal cost* component of the contribution is merely the next payment due on a payment amortization of today's present value difference between the AL and the PVB, or (AL-PVB); this is the mortgage payment version.[2] Or, alternatively, it is an amortization of the next payment needed, given today's present value for the AL, to build up to the future value of the PVB at some future time when retirement payments are to begin (the present-value-weighted average time to retirement, for the aggregate plan). This is the college savings account version. We'll use the former, because we can more readily obtain estimates for today's present values than for future values.

Using the HP-12C, or equivalent functions on a Microsoft Excel spreadsheet, we can readily calculate these amortizing payments on one sound basis, a method using level payments.[3] We'll need an interest rate and a time period. For the interest rate, let's do the analysis two ways, once using the actuarial expected return on assets (let's say 7.75%), and once using the risk-free long term government bond rate as recommended by economists and financiers (let's say 3.75%, even though this is a bit higher than actual rates, as of today).

For the time horizon, we need to make a fair guess for the average time remaining to retirement, when the debt is due. The actuary will likely calculate these normal cost and contribution values one employee at a time using the time remaining until probable retirement, and aggregate them. But a board member wanting to do a reasonableness check on whether the numbers have been "managed" unrealistically can't readily do that on his or her own; too many employees. So we want to guess at the average time until retirement across all the employees, and this should represent an average weighted by the present value of future benefit payments, PVB, for each one (this might be a WAG version of a

[2] I refer only to the next payment, as the amortization needs to be re-done for each period, given that asset values will have changed—even if the PVB were to be unchanged.

[3] There are many other methods in use by actuaries other than the level payment method. Level payments will usually give us a normal cost that is a bit greater than computed by most actuaries early in a plan's history, and a bit lower later in the plan's history. The normal cost estimate method outlined in the text ignores some or all of any expected growth in the liability, such as is related to cost of living increases, depending on what has already been included in our PVB and ABO input values; thus it may tend to understate normal cost.

guesstimate, although it could actually be refined with some actuarial help).[4] It will be shorter for an older and more mature workforce, and vice versa. Let's use 14 years as a starting point for this example, just to get started—and perhaps we'll be in the ballpark.

We'll do it first using the actuary's discount rate, the expected return on assets. So, on the calculator, clear all financial functions, enter 7.75% as the interest rate, 14 years as the time horizon, $300 billion (all figures are billions, but when inputting into the calculator drop the trailing zeroes just as I am doing, for convenience) as the actuary's value given for the AL, less the $400 billion given as the actuary's value for the PVB, or -$100 billion as the PV (mortgage payment version). We use zero for the FV. Hit PMT, and the answer, $12 billion, pops up. This, $12 billion per year, is the estimate of normal cost, and is the baseline portion of the contribution.[5] If using Excel, the formula is =pmt(.0775,14,300-400).

Contributions have to go up, however, when the plan is underfunded, and this gets us to situations analogous to what many plans are experiencing today. Let's say we have a funding shortfall, with only $200 billion in assets against the $300 billion AL. Ideally, this should all be paid immediately, but this is seldom or never possible given the highly volatile investment policies commonly used today and the financial situations of most sponsors. While like any past due debt, it should be paid immediately, there are various rules allowing deferral, addressing the length of the period over which this might be amortized. The good news is that these periods are (appropriately) getting shorter. But, for simplicity's sake, and being gentle and understanding of the employer's concerns about cash flow (and thus less sensitive to the employees' concerns about benefit security!), let's use the same horizon as for normal costs, the present value weighted average time to retirement, here 14 years. We use the same interest rate, the same horizon, and for present value we enter the value of the assets on hand less the assets that should be on hand, the AL, or $200 - $300 = -$100 billion. Because this difference in present values is also $100 billion, the makeup contribution answer ends up in this case being the same amount as in the normal cost portion, $12 billion. While the normal cost component should stay relatively stable over time if the plan is in steady state, only changing when the discount rate changes and as salary and benefit growth rates are

[4] See Waring (2011), Chapter 7, for more on refining this horizon estimate.

[5] The HP-12C calculates a payment that is level over time, and this is the type of payment used in most financing outside of pension plans. As mentioned, actuaries are accustomed to using a variety of other "shapes" for payments. One of the more common shapes, especially for pension plans, is a contribution growing at the same rate as wage growth, sometimes called "entry age" normal cost payments. A growing payment will start out a little lower than a level payment in the beginning, and be larger than a level payment in the end. It will also be a little lower on average across the aggregate of employees. The error introduced by using level payments is not very significant in most cases, as there are other errors in the reporting process that tend to understate the PVB, biasing the payment calculation in the opposite direction.

modified, the makeup contribution is dependent on the highly volatile asset portfolio's value and must be re-figured each period regardless of assumption changes.

The total contribution needed each year to restore the plan to fully funded status over the 14-year period is thus $24 billion, twice the baseline contribution and as a result perhaps quite difficult to meet.

We could have also gotten this total contribution, and in just one step, by simply amortizing the total PVB funding shortfall, skipping the intermediate accrued liability value. *I.e.*, amortize assets on hand, $200 billion, less the PVB, $400 billion, or -$200 billion, over the same 14-year period. This is the key to the highest level of simplicity for a board member wanting to do a quick reasonableness check on the adequacy of the current funding proposal: Amortize the assets less the PVB—that is, the PVB deficit (*not* the AL deficit)—over some guesstimated period of time that should approximate the present value weighted average time until retirement.

While these numbers are definitely approximations rather than precise calculations, they aren't bad approximations. And they are accessible to anyone that can run a financial calculator.

I suggested that we do this alternatively using the economist's assumptions, which are based on a risk-free discount rate reflecting the intention that the benefits be paid without risk of default. This is because the funding levels given in the opening paragraph of this article are based on *book* values for liabilities, not *market* values. If we use a market value approach, the AL and PVB values both appear much greater because of the lower discount rate, *although of course they represent exactly the same future benefit payment cash flows*. Figure 1 sets forth the key values and the required contribution components for both the expected return and risk-free rate discount methods side by side, and is loosely modeled after a particular large public plan:

Figure 1: Comparing a Large Public Plan's Finances Based on the Expected Return and on the Risk-Free Rate as their Discount Rate, Given a Current Level of Funding:

	Using Expected Return on Assets as the Discount Rate	Using the Long Term Gov't Bond Rate as the Discount Rate
Discount Rate	7.75%	3.75%
PVB	$ 400	$ 750
AL	$ 300	$ 550
Assets	$ 200	$ 200
Baseline Contribution (=Normal Cost), 14 year amort of PVB-AL	$12	$19
Support from assets: If assets deliver the expected return	N/A	($8)
Expected Baseline Contribution	$12	$11
Makeup Contrib, 14 year amortization of AL-Assets	$12	$33
Total Expected Contrib (Assets-PVB)	$24	$43

All figures are $ billions. Contributions are calculated on a level payment amortization basis, so that readers can duplicate any calculation made here, using an HP-12C ®. Duration for the baseline or normal cost contribution is set at 14 years, which is an assumption that the present-value-weighted average time to retirement for the members of this plan is 14 years. Makeup contributions are also amortized over 14 years, longer than some regulatory requirements and shorter than some past actuarial practices.

A couple of points are worth noting: If we include the expected return on assets when using the risk-free rate method, the expected baseline contribution (not the makeup contribution)—after considering expected investment return differences of 4% per year on the $200 billion of assets—is essentially the same or a bit less than for the traditional method based on the expected return of the assets. This is a mature plan, and the comparative levels of contributions for mature plans are discussed later in this article. The only difference is that the expected return on assets method *presumes* that the expected return is going to be achieved with certainty and incorporates that expectation through the discount rate, while the risk-free rate method doesn't count the returns until they are realized and so we have to put it down as an extra item to make sure that we are apples and apples.

The makeup contribution is significantly larger for the risk-free method than for the expected return on assets method. This is because we're working with the actual assets on hand now at the plan on which this example is modeled, not the assets that would have been on hand had they been using the ris-free rate method all along instead of the expected return method. So the AL deficit, on which the makeup contribution is based, is $350 billion instead of $100 billion, woefully short of a secure funding posture on a risk-free basis. So at this point in time, given the dramatic level of underfunding, such high contributions would be the price of making the benefit payments fully secure, and of stabilizing contribution risk.[6] After years of minimizing contributions, the public plan on which this example is loosely based would have a very difficult, perhaps impossible, task if it wanted to put its benefits on a true, risk-free, full-funding basis going forward.

In the following section we will contrast this first case, looking at what happens to contributions when both methods start out fully funded and then experience a comparable investment loss. In this case, expected contributions will be lower for the risk-free rate method than for the expected return on assets method.

Extended periods of disappointing market returns: Comparing contributions between the expected return on assets method of setting the discount rate, and the risk-free rate

If the risk-free rate accrued liability of $550 billion were fully funded with assets, there are a lot more assets working to protect the liabilities, and if invested in today's typical high-equity investment strategies the anticipated cost of the plan (conditional on actually realizing the expected return on average) will fully pay for the plan (see figure 2). The 4% per year additional expected return on the assets, above the risk-free discount rate, would give $22 billion per year of expected additional earnings that would help to offset pension expense and pension contributions—and in fact this is a bit more than the $19 billion of baseline contributions. So when the expected return is actually being realized, the plan would basically have zero contribution cost—a contribution holiday. So, when fully funded on a risk-free basis, the anticipated level of contributions is attractively low when using the risk-free rate as the discount rate.

The news is good, although not quite as good, if the plan is discounted at the expected return on assets, and also fully funded (again, refer to figure 2). Of course, if the invested assets actually do generate realized returns that, on average, are exactly equal to the expected return value of 7.75%, then the contributions will be exactly as estimated by the actuary and benefits will in fact turn out to be fully funded. The expected contribution is quite a bit higher for this plan than for the fund using the risk-free

[6] Note that I've been focusing on the comparison of the expected return on assets to the risk-free rate. This is a perfect comparison for public plans, and for GASB rulemaking discussions. Corporate plans, however, use for some purposes the corporate credit discount rate, which falls between the two extremes I've focused on. The inappropriateness of this rate is discussed more completely in Chapter 6 of my book (Waring 2011). It suffices to say here that the contribution results will generally fall between those from the risk-free rate method and those from the expected-return-on–assets method.

rate to discount its liabilities, because it has far fewer assets on hand yet faces the same ultimate benefit payment obligations.

But let's explore the implications of the fact that you don't always realize the expected return, and in fact sometimes there are extended periods of disappointing realizations. Markets are volatile, and they can deliver disappointing returns for long periods of time.

Let's consider first the plan using the expected return on assets. If in some one year, the assets return zero instead of the expected return of 7.75% (this is a not uncommon event, as it is less than a one standard deviation event where the investment policy has a typical 10% standard deviation), there is $300 billion times 7.75% of additional deficit, or $23.25 billion. This is a 23.25% increase in the PVB deficit, which was $100 billion before this loss. If the PVB deficit goes up by some percentage as a result of investment losses, the total contribution goes up by the same percentage (when the makeup contribution is amortized over the same period of time as the baseline contribution equal to normal cost). This means here that the makeup contribution adds 23.25% above baseline. That's uncomfortable, but perhaps survivable.

Figure 2: The Expected Return vs. The Risk-Free Rate: Comparing Contributions for a Plan Expecting to be Fully Funded, But That Experiences a Large Asset Shortfall After Disappointing Market Returns

If instead of the expected return planned in to asset growth, the markets only deliver a zero return for an extended period, the plan will be left with a large shortfall in assets--as has happened over the years 2000 through 2011. The question is, how much are contributions expected to increase in the following year, as a result? Holding investment policy the same:

	Discount = Expected Return on Assets		**Discount = Risk-Free Rate**	
Assumed Discount Rate:	7.75%		3.75%	
	Expected Position	***Actual Position After 40% Shortfall***	***Expected Position***	***Actual Position After 20% Shortfall***
Present Value of Benefits (PVB; always greater than AL)	$ 400		$ 750	
Initial Accrued Liability (AL)	$ 300		$ 550	
Assets	$ 300	$ 180	$ 550	$ 440
PVB Deficit (drives total contributions)	$ 100	$ 220	$ 200	$ 310
Baseline Contribution (=Normal Cost), 14 year amort of PVB-AL	$12	$12	$19	$19
Makeup Contrib, 14 year amortization of AL-Assets	$0	$14	$0	$10
Support From Assets: If Assets Deliver the Expected Return:	Included	Included	($22)	($18)
Total Contribution:	***$12***	***$26***	***($3)***	***$11***

Shortfall : As discussed in the text, after 12 years of essentially zero returns, and assuming some typical efforts to makeup for investment losses in a small way with additional contributions, we can expect a 40% or so asset shortfall relative to what they would have grown to had they earned the discount rate when it is based on the expected return. However, if the discount rate is the risk-free rate, that is the rate of return that is "baked in" to anticipated asset growth, not the expected return on the assets. So the shortfall is based on the loss of 3.75% per year for 12 years, rather than the loss of 7.75% for that period. Therefore, (and also assuming some modest makeup contributions) a 20% shortfall is all that we experience.
Support from the assets : The full expected return from the investment policy of 7.75% is presumed and included when the discount rate is the expected return on assets; no further correction to the anticipated contribution is needed. With respect to the case where we use the risk-free rate as the discount rate, however, and if the assets are invested in the same manner in both cases, we need to acknowledge that the difference in the expected return and the risk-free rate is anticipated to contribute to and reduce contributions.

But what if 12 years passed with an average return of just over *zero* on the assets, instead of a return equal to the expected return of 7.75% per year? (There is a reason to choose these numbers, a reason that I will get to in a moment.) A dollar of starting assets would have grown only to just barely over one dollar—but the actuarial expectation is that the assets would have earned 7.75% compounded for 12 years, growing a single dollar of assets to $2.45! So if we had made no makeup contributions during that 12-year period, the assets in the fund would be only 44% of what we had been counting on (assuming a stable plan)—a 56% shortfall!

On the other hand, if a makeup contribution in the amount of the investment shortfall would have been paid in to the fund each year, then the shortfall would never have grown to this level. But in fact, what is likely to have happened is somewhere between these two extremes: actuarial practice is to amortize the difference between actual and expected returns over many, many years, putting in makeup contributions but only slowly.[7] So, in practice, a portion of the difference might have been paid in—but the 56% shortfall is probably only reduced to (say) a 50% or 40% asset shortfall, still disastrous to the PVB deficit and thus disastrous to contributions.

Let's say we're on the better end of that range, and that despite 12 years of such poor returns we only have a 40% shortfall in assets, being short $120 billion of the $300 billion we planned on having. The PVB deficit (which controls contributions) will have grown from $400 billion minus $300 billion, or $100 billion, to $100 billion plus the $120 billion shortfall, or a total of $220 billion—a multiple of 2.2 times as large. All other things being equal, a payment on a debt is directly proportional to the size of the debt. And so, amortizing contribution payments would accordingly go up by the same multiple of 2.2, from 12 to 26. Overwhelming.

Now, consider the plan that uses the risk-free rate. This is much less sensitive. The assets only have to return 3.75% in this example in order not to disappoint, rather than 7.75% as in the expected return on assets method. That means that the shortfall after 12 years of the same flat returns is only about half as large for the risk-free rate method as for the expected return on assets method (also assuming that some makeup contributions have been made during those 12 years).

Also, this method doesn't count the expected return until it comes in as a realization. But if we are "expecting" returns—in the economist's sense—in the one case, we need to expect them in the other case in order to have a fair comparison. So we again add back in the value of the expected 4% premium of the expected return over the risk-free rate, to get the apples to apples comparison.

The bottom line is that while contributions for the plan managed using the risk-free rate do go up after this bad market event roughly by the same amount, $14 billion, as do the contributions under the expected return method, the total contribution under the risk-free rate method is still well less than the $26 billion required under the expected return method. Sponsors and participants would be having a better experience, lower longer term contributions and greater benefit security as a result of using the risk-free rate. More on why this is always going to be true as we continue this article.

[7] The expected return assumption is used when computing pension expense, rather than actual realizations, and accumulated realization errors are amortized back in only over very long periods of time. This is just one more method for deferring the cost of the pension plan. Such deferrals have a very high interest rate—the expected return on assets rate!

The reality of long term market volatility: Large-scale disappointing returns really can happen!

So, you say, such a period of extended disappointing returns couldn't happen: How could you have 12 years of near-zero returns on average? Don't the markets always deliver the expected return over long periods of time, so that the pension plan, being a long term investor, can count on getting the expected return?

Well, here's a news flash that is in front of all of us but that hasn't been much pointed out: From the end of 1999 through the end of 2011—the most recent 12 full calendar years—the S&P 500 has delivered just over zero percent per year, on average. To be precise, .55% —and that is with dividends included.[8] So the scenario that I have outlined is just exactly what happened.

One might guess, and appropriately, that this is one of the largest causes of today's pension crisis.[9] Since some plans made more makeup contributions than others, we see a good bit of variability, but the asset shortfall due to investments earning less than the expected return is the big item and the biggest reason most plans today are severely underfunded.[10]

"But we're counting on returns over the *long term*, Barton, and you're just showing a short little 12 year period." Well, that's a good question: Is 12 years long enough to be "the long term?" Interestingly, economists can't give an answer for how long the "long term" is! But as the question suggests, often people think of it as a longer time than 12 years, maybe 30 years or even 50 years. But there is no right answer.

But we might take a measure here from John Maynard Keynes, who is often quoted as saying "Markets can remain irrational a lot longer than you and I can remain solvent."[11] In light of this comment, I would suggest that if we're short 40%, plus or minus, of our expected asset value after only 12 years,

---

[8] Source: Morningstar (2011), and their Ibbotson Associates databases. By focusing on just the US stock market, the S&P 500, I'm of course exaggerating the problem a bit, as most pension plans are diversified also into other forms of equity as well as small portions of fixed income. Nonetheless, the exaggeration serves the point, as it is nearly universal across all plans that the expected return assumption has been greatly disappointed during this period.

[9] Another, of course, is the large benefit increases granted during the 1990s and early 2000s. But even so, plans were reasonably close to being fully funded at the beginning of this period.

[10] During the same period, long government bonds are up substantially, as interest rates have trended down. This means that the liability has increased in value substantially, another big "hit" to pensions, although one not fully acknowledged by the actuaries and the accountants. However, had plan sponsors hedged their pension plans by holding substantial amounts of this longer term debt, even if only with simplistic "LDI" solutions, they would have been protected from this rate change to that degree. Unfortunately, they didn't.

[11] Usually attributed to Keynes. This version from A. Gary Shilling, Forbes (1993) v. 151, iss. 4, pg. 236.

and experiencing sincere concern about our solvency and our ability to make up these shortfalls after that period, then perhaps 12 years does indeed qualify to represent whatever it is that we are trying to express when we talk about "the long term"!

Might we make up for this disappointing 12 year result with exceptional results in the next 12 years, justifying our faith that we will get the expected return over the long term? Yes, it is conceivable, but don't hold your breath. We would have to experience roughly twice the expected return of 7.75%, or about 15.46% over the next 12 years in order to make up the return shortfall purely on asset returns and without reliance on makeup contributions). What sober person expects this happy result, particularly given the current world economic environment?

Perhaps the remote likelihood of this eventuality can be highlighted by projecting the S&P 500 index out 12 years: If 2% of our needed 15% return were to come from dividends, so that the index itself needed to grow only at 13% in order to give us our 15% total, the index would have to grow from its year-end 2011 value of 1,259 to a value after the next 12 years of 5,457. While this would be a wonderful experience, any reliance on this remote possibility to save our pension plans is pure fantasy.

Why are contributions so volatile?

The key to understanding the extreme volatility of contributions is to note that a given percentage change in the value of the assets usually results in a much bigger percentage change in the value of the PVB deficit: If the PVB is $400 and the assets are equal to the AL of $300, a 10% swing in asset values, $30, becomes a 30% swing in the $100 PVB deficit—and thus a 30% swing in the contribution.

This volatility is obscured in practice for a while through many actuarial processes, such as smoothing, etc., which massage the true PVB deficit or the truly required contribution in the hope that improved investment returns will fix the problem without requiring the pain of makeup contributions.

But because returns can be disappointing for very long periods, this often doesn't work and in fact exacerbates the problem. Currently, for example, after 12 years of trying to avoid the pain it is now pretty clearly out in the open and the level of makeup contributions being acknowledged as necessary by plan actuaries has crept up substantially, but in the meantime funding levels have gotten disastrously low as a result of the slowly amortized response to market conditions. And of course, with the economy slow and both public tax revenues and corporate revenues slow, this is a difficult time to make any additional contributions whatsoever.

Shortfalls from disappointing investment returns should be made up immediately, in order to preserve full funding of the plan at all times. Such good habits are unlikely to develop any time soon—

actuaries and sponsors are firmly habituated to re-amortizing those shortfalls over quite long periods of time. I say re-amortizing—because the accrued liability is itself the product of an amortization. To then amortize asset shortfalls from the accrued liability over yet another large period of time is nothing if not the pure institutionalization of intentional underfunding. Two amortizations is, well, a lot of amortization. Would that my mortgage company were so generous with me.

These contribution volatility facts have not been widely understood. They are among the more important reasons to move to an economic or risk-free discount rate, where higher earlier contributions create a sufficiently large fund that contribution volatility is much dampened, and have the happy additional result that later contributions are smaller.

Do we always get the expected return over the long term?

It is perfectly ordinary for there to be long runs of good or bad market returns, as the markets tend to follow a random walk.

The use of the expected return as the discount rate is predicated on the false notion that a long term investor will get the expected return over the long term. This received wisdom is widely believed, especially by actuaries doing pension work. But no economist will agree with it; it is patently wrong; instead, the discount rate should be predicated on the market-related riskiness of the cash flows being discounted—*as it is in all other financing work.* For pension work, where the earned portion of the pension (the accrued liability) is intended to be secured through full funding and thus free of risk, this means a long term risk-free rate such as is observed on a long term government bond. But it is very difficult for actuaries and others used to looking to the assets for the discount rate to disassociate those assets from the discussion and look only to the liability itself.

The random walk theory has important implications contradicting the hope-based belief that we will get the expected return over the long term. Over the long term, a random walk means that the ups and downs of the asset portfolio's value—the *volatility* of the value of assets—will remain large and in fact will increase, not decrease, as the time horizon lengthens. It is important to absorb this fact fully and not to gloss over it. The expected return is just the center point of the distribution of possible future returns, and returns vary significantly enough that with compounding, the ending portfolio valuations will have a standard deviation that is a multiple of the one year standard deviation (more below). Over the "long term," however long that might be, the portfolio will on average go up in value with the expected return, but sometimes it will do better and sometimes it will do worse. And this might be much better—or much worse, as it has been during the last 12 years.

In fact, a portfolio will meet or beat the expected return (more precisely, the median equivalent of the expected return) only half the time over any particular long term time horizon, of whatever length—and the other half of the time it will fail to beat it. And as just said, it may beat it, or fail to beat it, by substantial margins.[12] No "magic" forces the value-added to average out at the expected return over long periods. It might happen—but it is more likely that that it will be higher, or lower, even much higher or much lower. Pension plans counting on the expected return thus have only a 50-50 chance (less, actually) of avoiding makeup contributions over any arbitrarily long period of time. And they have a very good chance that those makeup contributions will be bitterly large.

Take this seriously, as the disappointments can be huge: Assets are routinely invested in high-equity investment policies, with annual standard deviations of 10% and higher. The multi-year standard deviation of the ending portfolio value goes up roughly with the square root of the number of years elapsed. This means that over 16 years, the standard deviation is roughly 4 times the annual standard deviation, and over 25 years it is roughly 5 times.[13] This means that there can be quite long periods of disappointing performance, where asset shortfalls of the sort we have experienced over the last 12 years will happen. This is a normal event in the markets, one that has happened before and that we can assume safely will happen again. We should be planning for it.

Amortization of these shortfalls allows the total shortfall to become large when the period of disappointing returns grows long. Modest makeup contributions are made each year, based on the amortization, but each year the asset shortfall gets bigger by a multiple of the assets contributed. Today's low funding ratios reflect 12 years of near-zero asset returns, when returns over 7% and often 8% and even higher were "baked in" to expected asset growth, and only modest makeup contributions were amortized in.

It is worth remembering that under a random walk theory of the markets, we're always starting anew from wherever we are today, with no looking back. The expected return of 7.75% might be a fair estimate of the expected return over the long term—but it starts from today, meaning that we can't look back to a point 12 years ago and expect that history plus current history to give us the expected return. That's why it is called a random walk—if the expected return were to be expected over any long horizon from some past historical starting point, it would be a predictable walk, not a random walk.

[12] Because of the lognormality of accumulated returns, the probability of meeting or exceeding the *expected* return (rather than the *median* return) is somewhat less than 50-50. But 50-50 is close enough for this discussion . . .

[13] Again because of the lognormality of accumulated returns, these multiples are in reality a bit smaller on the downside, and a bit larger on the upside, than suggested by the square root rule. But the square root rule is close enough to teach the proper intuition for long term risk to the portfolio.

The right discount rate is theliability-matching or hedging rate, the risk-free rate, rather than the expected return

Let's spend another moment on the economist's method of determining the market value of the future benefit payments, discounting them using the risk-free rate when valuing the liability and when determining normal costs and thus also contributions: The reason that this discount rate is the correct one to use is that the stream of future benefit payments that we are valuing is intended to be free of risk. After all, it represents the deferred portion of pay, and employee's pay is considered sacrosanct and in almost all states has a special priority over other claims against the sponsor, a priority that should follow through to retirement payments. This intention that it should be free of risk is confirmed by the fact that it is required to be funded to the extent earned—that is, to the extent of the AL.

And, in the markets, future cash flows, whether assets or obligations, that are free of risk are discounted at the risk-free rate.

This is more than just an academic nicety—if we want to, we can hold bonds with cash flows matched to each future benefit payment cash flow, guaranteeing that each future benefit payment would be made when due.[14] Since this liability-matching investment structure would be free of investment risk, the value of that liability-matching asset portfolio is a natural and compelling basis for valuing the obligations it secures. On the other hand, if the liability is discounted at the expected return on the assets, there is no similar hedging portfolio available; it can't be done even if the sponsor wanted to do so! The liability is measured too small to be hedged, and it fluctuates with the returns of the markets rather than the interest rates that more truly impact its value.

There is no way to hedge all of the investment risk in the portfolio without viewing the plan in terms of the risk-free rate, inconsistent with the formal use of the expected return on assets. Thus the economists' and financiers' reliance on the risk-free rate to value the pension liabilities—only with it can a sponsor choose, if it wishes, to hedge all the market-related risks in the plan.

Over the long term, the contributions based on the risk-free rate are always less than those based on the expected return on assets, given the same investment policy

I suspect that the results shown above in figure 2 are surprising to most reading this article, that actual expected contributions for a plan are lower, not higher, when the discount rate is based on the risk-free rate. We can get some more intuition for this by conducting a mental experiment.

[14] In practice, this is slightly more complex and less precise for cash flows further off than about 30 years, or whatever the maturity of the longest-term risk-free bond is, but it is still "doable" in principle.

While not absolutely necessary, it is probably easiest to conduct this experiment by thinking in terms of a single employee just beginning her work life. We can more easily imagine a future stream of benefit payment cash flows starting at some future time of retirement, and an offsetting stream of contributions supplemented by investment returns between now and then (or even continuing on through the retirement years) for a single employee than for a thousand employees in aggregate. Yet one or a thousand, the mechanism is the same.

Regardless of the discount rate method used, the benefit payments in each year of retirement do not change; they aren't a function of the discount rate. And we also held investment policy the same in both cases in figure 2, so we'll also do that here. And we'll assume that the actual investments are heavily equity-oriented as is the custom today regardless of the discount rate method chosen. Thus the returns earned on the assets contributed and building up to eventually fund the benefit payments are the same in either case, both year-by-year and over the long term (just because we use the risk-free rate to value the liabilities does not mean that the assets in fact have to be invested in risk-free bonds; that is merely the lowest risk, lowest expected return choice available to the sponsor).

We then have two of the three key variables identical regardless of discount rate method used. The only variable between the methods, then, is the "shape" of the contribution payments over time—i.e., lower contributions earlier and higher contributions later, or higher contributions earlier and lower contributions later:

Using the expected return on assets as the discount rate, the starting contribution requirement (normal cost) is lower than the starting contribution calculated when using the risk-free rate. So the fund builds more slowly when using the expected return method, and even though it experiences the same *rate* of return the *dollars* of return will be less from that smaller fund.

By the same logic, the larger contributions from the risk-free rate method will build up a larger fund more quickly, and the larger fund will earn greater dollar returns from the same percentage investment returns.

Because the benefit payments are identical, then this must also mean that at some later point the relative size of the contributions will be reversed: The risk-free rate method will require smaller contributions later. And that is what we saw in figure 2.

If this isn't immediately clear, consider a *much* more streamlined, but still parallel, example: Let's say that we need to have a fund built up of $100 one year from today, and we'll make a payment today based on one of these same two discount rate methods. We'll invest the payment in the same high-equity investment policy no matter how big our payment is today.

Future Value in One Year: $100		***Final Payment or Refund, Based on a Realized Return on the Invested Initial Payment of:***			
Discount Rate Choices:	Initial Payment:	*1%*	*3.75%*	*7.75%*	*10%*
Risk-Free Rate: 3.75%	$96.39	$2.65	$0.00	($3.86)	($6.02)
Expected Return: 7.75%	$92.81	$6.26	$3.71	$0.00	($2.09)

If we make a larger payment (low discount rate) to that fund today instead of a smaller one (higher discount rate), then the makeup payment (or refund payment) required a year from now will necessarily be smaller (greater if it is a refund); i.e., you're always better off at the end if you put in more at the beginning. This will be true regardless of what the investment return is that happens to be realized, whether it is positive or negative or large or small. The table illustrates this improved ending status when using the lower discount rate, and points out that the amount of the makeup payment required (or refund earned) on one discount rate method might be a multiple of that for the other method.

This is completely analogous to the experience of a pension plan, even if simplified to the greatest degree possible. And if it were a pension plan, note that it is always better funded and the benefits more secure when contribution payments are made based on a lower discount rate rather than a higher one. Also note that it is only if the future obligations are discounted at the risk-free rate that there is a choice to hedge that obligation and de-risk the plan by investing in the risk-free asset—if the expected return on the assets is used as a discount rate, there can be no hedge (the assets will be insufficient to hedge the plan with long bonds, and an equity portfolio isn't a hedge). So, the end result must have a risky outcome, no safe outcome is possible.

Remember the start of this article? You don't want to get behind on the mortgage payment, but when things go bad in the investment markets the plans that have been using the expected return on assets as their discount rate will be seriously behind. And it will probably be difficult to make the contribution payments required in order to catch up.

And—as suggested by the analysis of the most recent 12 years' market history, above, we're in just such a period today.

Viewing the discount rate as a risk control dial

It is the exactly correct logic that the discount rate should be set by reference to the market-related risks in the benefit payment cash flows forming the liability, i.e., a risk-free rate, and not by reference to the expected return of the assets. But if you are one of the many that can't seem to disconnect the expected return of the assets from the discount rate for the liability and accept this logic, then please at least work from the agreed fact that the asset returns are risky, and they might—or might not—be sufficient to pay the benefit payments; they won't likely average out to the expected return, but will be above it, or below it: As demonstrated by the math of the random walk, you will be disappointed more than half the time if you're counting on the expected return to become a realized return.

And as demonstrated by our last 12 years experience, you don't get the expected return just because you are a long term investor.

So instead of referring to the intended level of risk in the benefit payment as the source of the proper discount rate, just think of the discount rate as a volume control knob for risk, a dial to control the pension funding risk, contribution risk, and benefit security risk of the plan. The higher the discount rate is set above the long term risk-free rate, the lower the sponsor's nearer term and intermediate term contributions will be. (But in the longer term they will be *higher* than if they had been based all along on the risk-free rate.)

If (it is really when) markets disappoint and fail to achieve the expected return for any substantial period of time—and they will slightly more than half the time if you use the expected return as the discount rate, but much much less often if you use the risk-free rate—additional makeup contributions will have to be made in order to make up for the missing investment earnings.

When you turn up the risk control volume control, funding levels are lower, future contribution requirements will be greater, and in disappointing markets the future contribution requirements will be much greater. Because makeup contributions are only needed when the market has been down for extended periods, and because they can be very large even when amortized over very long periods, this discount rate choice is also a dial controlling the security of benefits. The higher the discount rate, the lower the baseline (normal cost) contributions are but the higher the surprise contributions will be when markets disappoint, and the lower will be the security of benefits.

The most predictable and safest discount rate—for the sponsor and for the beneficiaries—is the method using the risk-free rate. If you want to take all the risk out, you can even hedge this measure of the liability; in this case you have turned on the risk control "mute switch."

But if you use a higher rate, just how loud do you want the risk noise to be? Maybe it is better to ask how loud can you afford the risk noise to be—over the long term? Although most sponsors have chosen the expected return as the discount rate in the past, it is my expectation that with a full understanding of the risk involved, most sponsors (and most employee organizations) would prefer a lower risk level rather than a higher one. They like predictability and they like safety. What sponsor or plan participant feels good about the pension funding situation we find ourselves in today?

The discount rate argument is, in effect, just a proxy for an argument about how large the contribution should be in the early period of the plan or when a plan is underfunded. And this is equivalent to an effort to justify deferring some large portion of the contribution from today to later. (And deferring it at a higher rate of interest, which is what the discount rate really represents.) But sooner or later the benefit payments have to be made, which are the same regardless of the discount rate—so the two methods sooner or later have to tie out, either more contributions early and less later, or less early and more later.

If you do use the risk-free rate to discount the benefit payments, there is nothing to keep you from using an investment policy containing a lot of equities. As shown in figure 2, this might well provide substantial help in paying for the plan in the long term, but it does add risk that it will make the plan more expensive if and when markets disappoint and fail to earn an average rate equal to the risk-free rate. It is a healthy thing that you won't be counting the returns and using them to lower your future contributions when you form your expected return assumption, but only as those returns are *realized*. Count your chickens when they hatch, not before.

How do we get there from here?

It is a regrettable fact, but a fact nonetheless, that safe and secure pension benefits are more expensive than we have been told, because the expected return method of setting the discount rate implies a high level of previously undisclosed benefit security risk. Risk that we have now experienced.

Perhaps sponsors can find a way to increase their contribution levels. But it is likely that labor representatives will have to reassess the level of retirement benefits that they deem reasonable if they want those benefits to be secure. Yes, there are huge and difficult political, legal, and contracting problems in modifying labor's benefit expectations, especially for benefits already accrued. But we're seeing examples now—more so outside the US than inside, but it is happening here also.

But we have seen the bankruptcy courts used as a blunt object to simply get rid of these benefit obligations at some corporate plans, and it is being tried tentatively now for some smaller public plans. Might it not be better to use these methods to modify benefits in a more refined way, to remake them in a sustainable fashion rather than to get rid of them?

I have said often in lectures, and also in one article and in one book chapter, that the worst DB retirement plan benefit is far better than the best defined contribution retirement plan benefit.[15] DC contribution rate are limited at low levels; participation is optional, investments are expensive and

[15] Waring (2007) and (2009).

investment policies are poorly developed, and there is no inexpensive means to annuitize the balance. The sad result is that median balances in DC plans at time of retirement are in the order of magnitude of $100,000 and often less. The stingiest DB plan will provide a much better retirement income than will such a balance.

So—it is imperative to make DB plans work, and that will take an open mind to reassessing the cost of safe and secure benefits and that will necessarily require a reassessment of what a "good" DB benefit is.

With that, if contributions are calculated on a risk-free basis, and if fully funded and kept fully funded then these new DB benefits can be very secure indeed, now and for generations to come.

Conclusion

There really isn't any way out of this argument against the use of the expected return as the discount rate, and in favor of the liability hedging rate—the risk-free rate—regardless of the investment policy chosen. It is unequivocally true that those who desire their DB plans to be healthy, securely funding their promised benefits, need to be advocating the use of the risk-free rate in place of the old-school actuarial expected rate of return. It was perhaps a good-hearted effort by the actuarial community to save us all money in funding DB plans, but the savings from their expected return discount rate are not only illusory but reverse later; so it was a badly misguided effort, neglecting the effects of investment risk. Those low early contributions create a very real and very large risk to the solvency of the plan over any reasonably long period of time (after a "long term" of only 12 years, right now!),[16]

The risk-free rate is the rate to use, if only for management purposes rather than accounting purposes, whether you are regulated by GASB, FASB, or IASB. The fine points of whether the wrong rate in use is the expected return or some other rate such as the high quality corporate credit rate is just that, a fine point: It is the wrong rate, the volume control knob for risk is turned up too high, and it will to that degree be exposed to larger risk of underfunding, of unhappy contribution surprises, and benefit insecurity.

[16] One of the very sad ironies to note here is that the 30 year returns of the S&P 500 are well above the expected return assumption—even though they include the disastrous most recent 12 years—averaging over 11% per annum. This is because the two decades of the 1980s and 1990s gave us market returns far *above* expectation. But alas—these gains were given away in the form of benefit increases, rather than being saved against the future rainy day. A rainy day that in fact has since happened.

You can pay later, or you can pay earlier, and the discount rate is the determinant. But what you pay works out the same in present value terms, regardless. Not surprisingly, there is no free lunch from a high discount rate.

References

Morningstar Inc. 2011. Ibbotson Stocks, Bonds, Bills, and Inflation (SBBI) 2011 Classic Yearbook. Chicago: Morningstar, Inc.

Waring, M. Barton, and Laurence B. Siegel. 2007. "Don't Kill the Golden Goose! Saving Pension Plans." Financial Analysts Journal,vol. 63, no. 1 (January/February): 31-45.

Waring, M. Barton. 2009. "Between Scylla and Charybdis: Improving the Cost Effectiveness of Public Pension Retirement Plans." In *The Future of Pubic Employee Retirement Systems*. Edited by Olivia S. Mitchell and Gary Anderson. New York: Oxford University Press.

Waring, M. Barton. 2011. Pension Finance: Putting the Risks and Costs of Defined Benefit Plans Back Under Your Control. Hoboken, NJ: John Wiley & Sons, Inc.