CENTER ON
FEDERAL FINANCIAL
INSTITUTIONS

Suite 600
2121 K Street, Northwest
Washington, DC 20037

TEL: 202.775.6329
WWW.COFFI.ORG

original issue date: September 10, 2004
Douglas J. Elliott
douglas.elliott@coffi.org

PBGC: A Yield Curve Primer

Various experts, including the Department of Treasury, advocate the use of a "yield curve" to improve the measurement of the liabilities and funding requirements for a "defined benefit" pension plan. This primer attempts to explain the concept in non-technical language. Please refer to www.coffi.org for other related papers on pension promises and the government corporation that guarantees them, the Pension Benefit Guaranty Corporation (PBGC). "PBGC: A Primer" would be particularly helpful for those unfamiliar with pension plans and their funding.

From a theoretical point of view, a yield curve approach increases accuracy of liability measurement. Opposition centers on whether the gain in theoretical accuracy is large enough to offset such negatives as increased volatility of funding requirements.

We address the following questions:

- What is the yield curve proposal?
- Why should the discount rate vary?
- Why do bond prices rise when interest rates fall?
- Why would a yield curve approach be better?
- Why do different pension funds have different payout patterns?
- Why would anyone oppose a yield curve approach, if it is more accurate?
- How would the proposals affect "smoothing"?
- Would using a "spot" rate be better?
- How do financial economists view this question?
- Are the financial markets efficient enough to warrant this technique?
- What would the practical effect be of using spot rates?

What is the yield curve proposal?

Federal law requires that companies make contributions to ensure that the value of assets in pension funds that they sponsor roughly equals the value of the pension payments that have been promised. Since the pension payments will be spread over many years, it is necessary to calculate what those payments would be worth in today's dollars. (For many reasons, people prefer money now to the same amount of money in the future.) This calculation involves reducing the value of future payments based on (1) the number of years until payment and (2) an interest rate, called a "discount rate." (See the appendix under "What is a Discount Rate?")

Current law uses the same discount rate for all time periods. Yield curve proposals aim to improve accuracy by using an appropriate set of rates, one for each year in the future. The set of rates is called a yield curve, because they are commonly shown as a graph with a curving line depicting the rate appropriate to each length of time. Yield curves have a long history as a tool used by the financial markets to analyze interest rates.

Figure 1, below, illustrates the yield curve for Treasury notes and bonds, as of June 2, 2004. Note that the curve generally slopes upward from left to right, which is referred to as a "normal yield curve," since it is the most common shape.

Figure 1

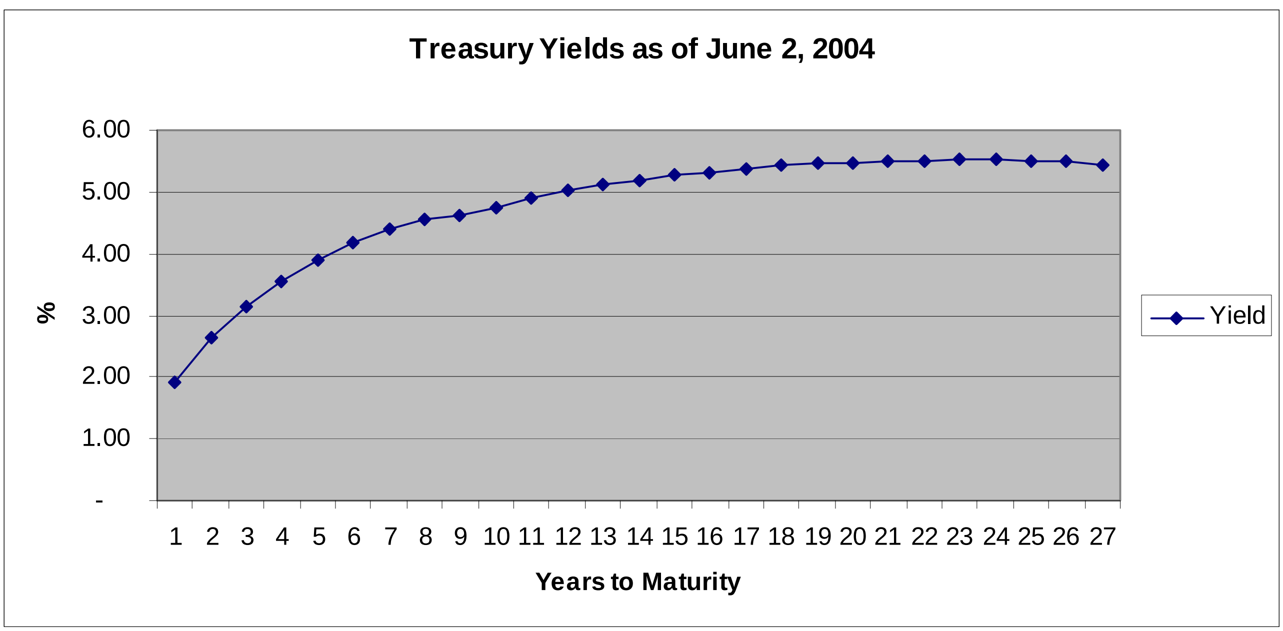

Source: Yields from *Wall Street Journal*

We have illustrated the yield curve using rates on Treasury bonds, but a yield curve can be calculated for any type of fixed income instrument where the interest rate varies with maturity. In particular, the Treasury Department's proposal uses a yield curve for high-quality corporate bonds.

Why should the discount rate vary?

There is disagreement among experts about how to determine the right discount rate for pensions (see Appendix). However, there is a strong consensus that the rate ought to reflect interest rates available in the financial markets. For example, the history of previous pension legislation suggests that Treasury rates were chosen originally, at least in part, because they were seen as a reasonable proxy for the rates insurers charged for group annuities. (A group annuity is a contract whereby an insurer takes over the obligation to make pension payments to a group of present and future retirees.) The recent move to using a discount rate based on corporate bonds (supported by both Congress and the Administration) had as a prime rationale the argument that such a rate better reflected pricing on group annuities.

For all types of bonds, financial markets consistently provide different returns for different maturities of investments. For example, the U.S. Treasury found, as of June 2nd, that it must promise a 4.73% return on money that it borrows for 10 years, but needs only to promise 1.92% for 1-year borrowings.

Generally, investors will demand a higher interest rate for longer-term borrowings, essentially charging for giving up the flexibility to move their money to another investment. Borrowers, for their part, are willing to pay more for the stability provided by longer-term funding. This effect is more pronounced for less credit-worthy borrowers, since investors are more comfortable estimating bankruptcy risk in the near-term than in the long-term. Thus, Treasury bonds tend to show a less steep yield curve than corporate bonds, since the U.S. government is viewed as having virtually no bankruptcy risk.

However, there are some economic conditions under which long-term money earns a lower return. An impending recession can produce this effect, if investors perceive that interest rates will fall as future economic activity, and inflation, slow. Some investors would choose to lock in a lower rate for a longer period rather than temporarily receive a higher rate that would be offset by lower rates when the funds are reinvested. These circumstances are said to create an “inverted yield curve,” since the graph slopes down instead of the more usual upward slope.

Figure 2

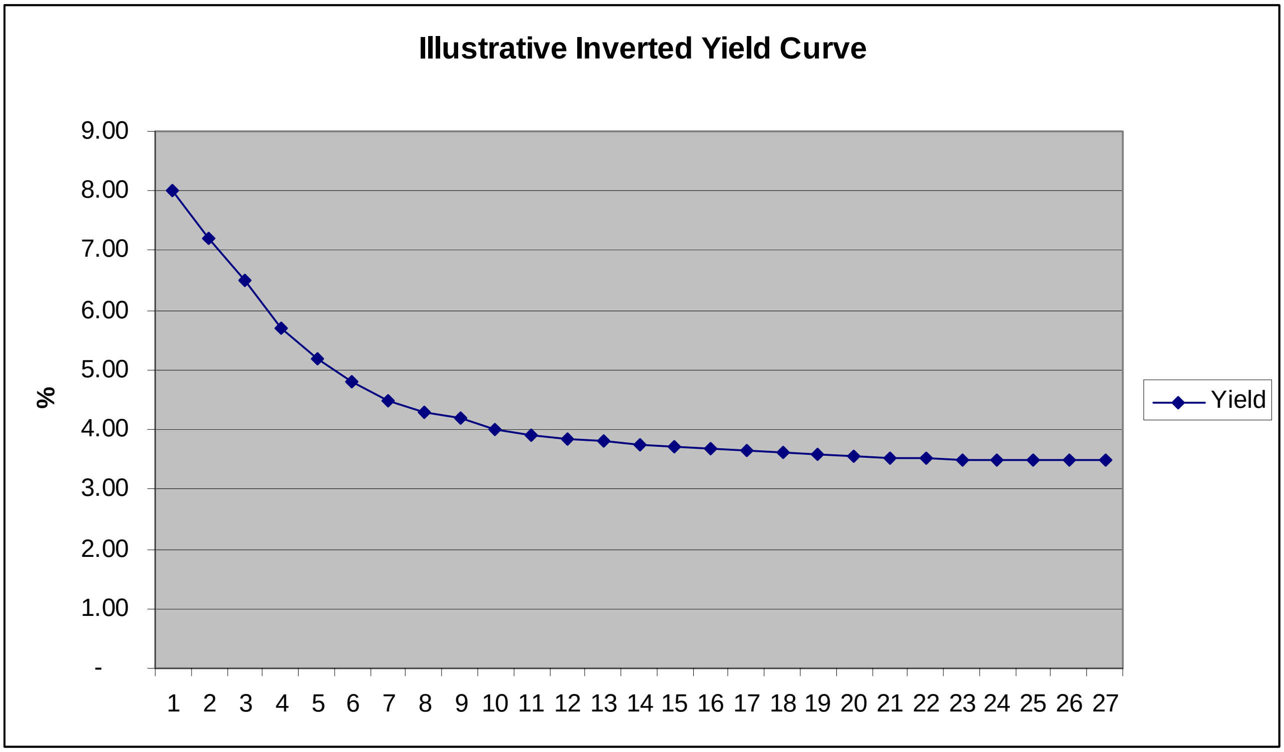

Source: COFFI estimates, for illustration only

Inverted yield curves have always been substantially less common than normal yield curves and have been rare in the last 20 years. Peter Fisher, then Undersecretary of Treasury, testified on July 15, 2003 that "over the last 20 years, the Treasury yield curve has inverted a total of only 14 months out of 20 years. The corporate yield curve, which we are recommending as the basis for the yield curve for measuring liabilities, was inverted for only 1 month out of the last 20 years, and, at that, only a very fraction of a few basis points [sic, in transcript]." Much more common are shifts in the degree of tilt in the normal yield curve. These can be as significant in their effects as an actual inversion.

For completeness, we should note that a yield curve does not need to be steadily increasing or steadily falling across the maturity spectrum. There are occasions when a kink develops, such as after 30-year Treasury bond prices rose sharply (reducing the market interest rate) after the Treasury Department announced it would cease selling new 30-year bonds. This created a scarcity premium for the longest maturity bonds.

Why do bond prices rise when interest rates fall?

Many observers find it counterintuitive that bond prices and interest rates move in the opposite direction. The reason is actually a simple one. If a bond is sold for $100 at a time when the market demands a 5% interest rate, then the issuer will promise a $5 interest payment each year (5% times $100.) If market interest rates have risen to 6% at the point where the bondholder wishes to sell, then he or she will have to lower the price in order to find a buyer. Otherwise, potential buyers will always prefer to exchange their $100 for a new bond paying $6 a year. The amount of price reduction will depend on the time until the principal will be repaid (the maturity). For example, if there is only one year left to go, then it may only take a $1 price reduction, to $99, to compensate for the difference of $1 in interest payment per year between the old bond and new ones available in the market. If there are 20 years left to go, the price reduction will have to be closer to $20.

Why would a yield curve approach be better?

A set of discount rates based on a yield curve should provide a more accurate measurement of funding requirements than a calculation using a single rate. After all, pension funds make payments each year, not just one payment at an average point in the future. (In theory, a quarterly or monthly yield curve would be even more accurate, but no one has proposed that such precision would be worth the effort.)

Current law provides for a discount rate based on the average yield of long-term corporate bonds of high credit quality. This should be fairly accurate for a pension fund whose payments are likely to be similar to the schedule of principal and interest payments of a bond of the chosen maturity. However, problems can arise when a fund's expected payments differ significantly in maturity from the underlying bonds used to calculate the single discount rate.

For example, a fund with a much shorter time until its average pension promise comes due may find that the rates available for that shorter period are significantly lower than the discount rate. This would mean that holding investments equal to the present value calculated from that discount rate would produce funds insufficient to make all the payments. The fund would appear to be fully funded, but would actually have inadequate assets, unless investments were made in higher-earning assets, with their attendant higher risk of failing to earn as much as needed.

The opposite problem is possible for a fund with a much longer time until its average pension payment is due. The single discount rate may be lower than available market returns, forcing more contributions early on than will be required to make the eventual payments. This overfunding improves benefit security but discourages employers from providing pensions.

Group annuities provide the closest market parallel to pensions. Proponents of the yield curve approach point out that insurers consistently use yield curves when pricing group annuities.

Why do different pension funds have different payout patterns?

To paraphrase Peter Fisher's testimony:

> Each pension fund has a unique schedule of future benefit payments – or cash flow profile – that depends on the characteristics of the work force covered by the plan. These characteristics include:
>
> - The percent of participants that are retired
> - The age of current workers covered by the plan
> - The percent receiving lump sums
> - Whether the covered work force has been growing or shrinking over time

> Plans with more retirees and older workers, more lump sum payments, and shrinking workforces will make a higher percentage of their pension payments in the near future, while plans with younger workers, fewer retirees, fewer lump sums, and growing workforces will make a higher percentage of payments in later years.

Traditional manufacturing industries, a key part of the defined benefit universe, tend to have shorter payouts.

Why would anyone oppose a yield curve approach, if it is more accurate?

Opponents generally do not argue against the theoretical merits, but they believe that the search for purity comes at too high a cost. The arguments fall into three categories:

An additional source of volatility. The funding requirements faced by companies currently depend, in part, on changes in the average interest rate on long-term, high-quality corporate bonds. A yield curve approach extends this to a series of interest rates for bonds of all maturities from 1 year to 30 years (or further, if the Treasury Department analyzes further out). Historically, interest rates are more volatile the shorter the maturity. (Note, for example, that the Fed Funds rate, a daily rate essentially controlled by the Federal Reserve Board, dropped from 6.5% in late 2001 to a low of 1.0% in mid-2003, while the 20-year Treasury Bond rate fell only about a point and a half, from roughly 5.75% to 4.25%.)

Volatility is also increased by changes in the demographics of the workforces. The maturity of the pension promises can be changed by such factors as layoffs, early retirement programs, divestitures, acquisitions, and changes in actuarial estimates of lifespans. A change in the timing of payments would change the discount rate under a yield curve approach but would not under the current approach using a single discount rate. The discount rate change in turn would affect funding requirements, increasing them as discount rates fell or decreasing them as they rose.

Additionally, matching assets to liabilities can become more difficult. Some firms try to stabilize their pension funding requirements by investing in bonds with maturities similar to that of the index used to calculate the discount rate for legal pension funding requirements. Such a company would now need to take account of possible changes in the slope of the yield curve. That is, the company would need to consider not only changes in the general level of long-term rates, but changes in the relationship of long-term and short-term interest rates. Of course, most yield curve advocates would argue that this greater precision in matching is desirable.

Only a minor increase in accuracy. Many opponents believe that the detailed analysis by the Treasury Department, when it comes out, will show that there will be relatively little difference between the discount rates calculated on the current method versus the more refined approach. Given the long life-spans in retirement, even pension funds with a high proportion of retirees still have quite a long average maturity. It is unclear that there has been a significant difference in interest rates historically between say, 15-year and 25-year interest rates. (The difference, for Treasury securities, was 0.23%, as of June 2, 2004). Part of the debate will be about how large a difference is significant for public policy purposes. This argument may resolve itself once the Treasury Department provides its detailed proposals and the analysis behind them.

Implementation difficulties. Some opponents believe that it will be difficult in practice to calculate accurate yield curves. One potential problem is that there are few high-quality corporate bonds at certain maturities, creating the possibility that a significant improvement or decline in the creditworthiness of a single company or a particular industry could unduly change the discount rate. A second issue is that some future pension payments are of longer maturity than all, or virtually all, bonds. There are technical "fixes" available for both problems, but they require making additional assumptions that detract from theoretical purety and are subject to disagreement. Yield curve proponents believe these approximations are unimportant compared to the gains in accuracy.

How would the proposals affect "smoothing"?

Under current law, the discount rate used for pension funding requirements is based on an average of the market rates of three years ago, two years ago, last year, and now. The calculation uses a weighted average to give more weight to recent data points. 40% of the weight is given to current rates, 30% to last year's, 20% to two years ago, and 10% to three years ago.

There is no theoretical reason why a weighted average yield curve could not be calculated. There may be a transition issue, if the Treasury Department finds that information not previously collected must be gathered going forward. However, this problem would vanish within three years as more complete information entered the historical record.

In practice, advocates of the yield curve approach generally propose using only current rates (called "spot" rates) in the calculations. This separate change could have significantly more effect on the discount rate than the use of a yield curve itself, depending on the speed with which interest rates change in the future.

Would using a "spot" rate be better?

Using spot rates to calculate the discount rate has the advantage of immediately reflecting present market conditions in determining whether a pension fund is capable of meeting its long-term obligations. It has the disadvantage of increasing the volatility of funding requirements, since spot rates fluctuate more than historical averages.

Use of a spot rate is particularly beneficial when the financial markets make a rapid, but lasting, adjustment to changing conditions. For example, if the recent bursting of the bubble signifies a long-term return to more normal conditions, then it likely is best for accounting and regulatory measurements of all kinds to reflect this quickly. A significant unexpected event, such as an outbreak of war, is another example of a situation where historical averages could become misleading guides to the future.

On the other hand, if market fluctuations are part of a normal cycle, and will be reversed, then smoothing avoids an over-reaction to a short-term factor. For instance, interest rates tend to fall in recessions, which would drive up funding requirements more sharply under a spot rate approach. This is unfortunate, since a recession is the toughest time for companies to fund their plans. Some argue that a normal business cycle brings interest increases and decreases that will average out over years and that, therefore, longer-term averages should be used to calculate discount rates.

How do financial economists view this question?

Financial economists generally support the use of spot rates, based in part on a belief that the U.S. financial markets are relatively "efficient". If financial markets were perfectly efficient, then the spot price of a stock or the spot interest rate of a bond would be the best possible reflection of its true value, based on all information currently available. Only new information would cause rate movements to differ from those now predicted and we could not know in advance in which direction new information would move the market. Financial theory assumes that arbitrage activity will push prices back into line with the best information available. For example, if interest rates temporarily fall too far, then smart holders of those bonds could benefit by selling their bonds while they are overpriced. These sales would push prices back down and interest rates back up.

The yield curve contains implicit predictions about future interest rate movements. For instance, an investor who has money to invest in bonds for two years could either buy a two-year bond or could buy a one-year bond and plan to reinvest a year from now. Currently, markets expect short-term rates to go up over the next year, so the two-year bond has a higher interest rate than the one-year bond. Otherwise, if they had the same interest rate, investors would buy the one-year bond now, figuring on reinvesting in a year at the higher rates, thus earning a larger total return for the two-year period. Investors would therefore bid up the one-year bond compared to the two-year, pushing its interest rate down.

If current interest rates reflect the best available information about future interest rate movements, then policymakers will be taking a gamble in using historical averages. Unfortunately, the gamble would be akin to playing roulette at a casino. Sometimes the historical averages will be a better predictor of future interest rates than would be a yield curve using spot rates. However, most of the time, the market predictions would win out, assuming perfect efficiency.

Are the financial markets efficient enough to warrant this technique?

Many opponents of the use of spot rates are not fully convinced by arguments of market efficiency. The same general efficiency argument applies to stock prices and we have just witnessed a great run-up in stock prices, particularly for technology stocks, followed by a sharp decline. There is a common perception that financial markets tend to overshoot in both directions, which intuitively argues that historical averages be used.

A more technical point is that short-term interest rates are not completely market-determined, since the Federal Reserve Board uses these rates as an instrument of public policy. The presence of a massive investor in very short-term bonds with non-economic objectives falls outside traditional arguments of market efficiency. However, the Federal Reserve has considerably less effect on longer-term rates, such as those that principally impact pension funds.

In the pension arena, the arguments about the "perfect storm" that has struck pension funds usually make two assumptions. First, that interest rates overshot when they fell and, second, that they will revert to levels more similar to those of the 1980's and much of the 1990's, when rates were considerably higher. (As is shown in "PBGC: Fundamental Questions," available at www.coffi.org, rates on long-term bonds are already consistent with a reversion to historical averages for the period since the end of World War II.)

What would the practical effect be of using spot rates?

If everything else were equal, funding requirements would bounce around more. Funding needs would decline when interest rates rose and increase when rates fell. Eliminating the smoothing provided by historical averages would necessarily increase the volatility of the discount rates used for pension funding rules. However, the degree of the effect would depend heavily on the speed with which rates.

This greater movement can be viewed as positive, if it is seen as a better, quicker adjustment to new information, or negative, if it merely slavishly follows short-term market moves. Note that some yield curve proponents have suggested that other changes in funding rules could be used to offset the resulting increase in volatility, while retaining the benefits of more accurate measurement.

Appendix

Reprinted below are several sections of "PBGC: A Primer" that may be particularly helpful for those unfamiliar with the details of pension policy. The full report is available for downloading from www.coffi.org.

What is a discount rate?

There are many times when one needs to compare payments or receipts that fall into different time periods. For example, how does a lump sum pension payment of $10,000 taken now compare to receiving $100 a month for an expected life of 25 years?

Economists, actuaries, accountants, and financial markets generally calculate "present values" to make these comparisons. Each payment or receipt is "discounted" to an equivalent value in the present moment by using (1) a factor based on the number of years before money is exchanged and (2) an interest rate called a "discount rate." Decisions can then be based on a comparison of the present value of the payments and of the receipts.

For example, a payment of $1,000 in 2 years, discounted at 5%, is approximately $907. It tends to be easier to understand this by reversing the process. If we put $907 today into a savings account earning 5% per year, it will grow to a value of $1,000 in 2 years. That is, $907 earns $45 in interest at 5% and totals $952 at the end of year 1. $952 earns $48 in interest and equals 1,000 at the end of year 2. The mathematics can become complicated, but the logic always remains that of a present value growing at an interest rate for some number of periods to equal the ultimate target amount.

How do discount rates affect pension calculations?

Present value calculations, and therefore discount rates, are critical to pension calculations, since the benefit obligations are spread over many years. There are three principal areas where present value calculations affect pension plans:

Funding. The adequacy of funding in a pension plan is determined by comparing the present value of the pension obligations to the value of the pension investments. These funding levels help determine the legally required contributions from the plan sponsor and the variable premiums that PBGC charges some plans. A higher discount rate means that the present value of pension obligations is correspondingly lower, and therefore lower contributions are required to fully fund benefits.

Lump sums. The lump sum payment offered to certain employees as they retire represents the present value of the future benefits they would otherwise receive. Higher discount rates produce lower lump sums.

Terminations. In practice, with rare exceptions, PBGC will not take over a pension trust unless the present value of future benefit payments, calculated under PBGC rules, exceeds the value of trust assets. Otherwise, the pension trust should be able to pay its own way, even without further contributions from a plan sponsor.

How should discount rates be determined for pensions?

There is a vigorous debate on the correct discount rate to use to calculate the adequacy of pension funding. This debate is far from academic. Using a one percentage point lower discount rate for funding purposes can increase the calculated liability by 15%. At the considerable risk of over-simplification, this debate can be summarized as follows.

Expected rates of return. Actuarial rules specify that the discount rate should be based on the expected long-term return on investment in the pension trust. Thus, a plan would be viewed as fully funded if it has investments sufficient to make the required payments assuming the assets do indeed earn the estimated rate. For some years now, actuaries working for plan sponsors have estimated the return for the average plan at about 8%, although estimates for individual plans have varied up or down by one or two percentage points. Accountants also use the expected investment return as part of the calculations of funding levels under the relevant Financial Accounting Standards (FAS). (The recently promulgated FAS 132 expands on FAS 87 and FAS 17, which constituted the principal rules for pension accounting for private sector companies.)

Risk-free rates. Financial economists generally advocate the use of a “risk-free” rate, such as the return on a Treasury bond of equivalent duration to the obligation. Currently, this would produce a discount rate of 4% to 5%. This might at first seem to conflict with the well-established use of higher discount rates to value riskier investments, so we will start by explaining that methodology.

People and institutions are assumed to dislike risk and to demand a higher return on riskier payment streams. This is reflected by choosing a higher discount rate. Thus, a stock expected to earn 8% a year might be viewed as fairly valued when an 8% discount rate is appropriate for the risk, but would be under-valued if the risk warrants a discount rate of only 7%. Generally, bonds will be less risky than stocks, because they have stronger legal claim on payments from the issuing company. Therefore, they will have lower discount rates.

However, these discount rates are all from the point of view of the recipient of the future payment, who will demand a higher expected return to compensate for the risk of not being paid. There is no such uncertainty for the one who promised the payment, assuming the promises are clear legal obligations, and therefore no reason to increase the discount rate. From the payer’s point of view the payment is just as certain as a Treasury bond payment is from the recipient’s point of view.

Some argue that the ability to pass the pension obligation to PBGC renders this analogy irrelevant for pensions, however this applies only in extreme circumstances. First, corporations try very hard to avoid insolvency, which is a necessary practical condition for PBGC help. Managements seldom treat future promises as less costly because they could theoretically be evaded in bankruptcy. Second, this argument would apply to virtually any contractual obligation, since bankruptcy potentially affects all claims. Yet no one is suggesting that a high discount rate be applied to all future payments.

Rates on high-quality bonds. There are at least three justifications for using high-quality bond rates. First, it represents a compromise between traditional actuarial practice and the arguments of the financial economists. High-quality bonds are not truly risk-free, but they do tend to have a low probability of default. Second, the implied interest rates that insurers build into their pricing when they offer to take on pension obligations are generally close to high-quality bond rates, if one factors out the additional administrative and other charges that are also built in. Third, it is a convenient political compromise between the arguments of different groups that favor risk-free rates or rates closer to long-term investment returns.

What discount rates are used for legal purposes?

Congress has not settled on a single intellectual position in its various applications of discount rates to pensions. The discount rate rules for pension funding are currently in flux. H.R. 3108, if signed into law, would temporarily change the rules for 2004 and 2005. Congress intends to enact permanent changes of some nature before these provisions expire.

For the next two years, plan sponsors may calculate current liabilities by choosing a discount rate in a “permissible range” from 90% to 100% of a four-year weighted average of high-quality corporate bond rates. The Treasury department will compile an index based on two or more publicly available indices that take account of the rates on thousands of bonds traded in the markets. The four-year weighting of these indexed rates is likely to be designed to give considerably more emphasis to recent rates.

The funding rule immediately prior to the new Act used an average of rates on the 30-year Treasury bond and a “permissible range” of 90% to 120%, but was otherwise identical in construction. Corporate bond rates will be higher, and pension funding obligations lower, since investors charge a greater yield to account for the added risk of bankruptcy.

PBGC assesses an additional “variable rate” premium for certain underfunded plans (see “How is PBGC financed?”) It calculates the level of underfunding using a 30-year Treasury rate. Currently, it uses 85% of that rate, down from 100% used in 2002 and 2003. The proposed Act retains the 85% ratio, but applies it to a corporate bond rate.

Lump sum payments to participants are generally calculated using the 30-year Treasury rate, the highest rate currently allowed. Benefits paid as lump sums are considerably higher than the present value of those same benefits paid as monthly payments in retirement, as a result of a gap between discount rates used for lump sum and funding calculations. The proposed Act increases this gap by raising the discount rate for funding purposes without altering the lump sum discount rate. There is little theoretical economic justification for using different rates.