MILLIMAN WHITE PAPER

# Constructing discount curves under LDTI

Josh Dobiac, LLM, JD, MS, CAIA
Bill Matczak, FSA, MAAA
Jeff Greco, FRM

Milliman

The Financial Accounting Standards Board (FASB) issued new accounting standards on August 15, 2018, containing targeted changes for how long-duration insurance contracts are accounted. Initially required for financial reporting after January 1, 2021, this standard has been delayed until January 1, 2022, for large companies and January 1, 2024, for small and medium-sized firms. The new standard addresses four key issues:

1. **Improve the timeliness** of recognizing changes in the liability for policy benefits of nonparticipating traditional and limited payment contracts and modify the rate used to discount future cash flows.
2. **Simplify the amortization** of deferred acquisition costs.
3. **Simplify and improve the accounting** of certain market-based options and guaranteed benefits associated with account value-based contracts.
4. **Improve the effectiveness** of required disclosures.

GAAP Long Duration Targeted Improvements (LDTI) has made substantial changes to the valuation of insurance liabilities. These changes are particularly pronounced for nonparticipating traditional liabilities. For these products, liabilities must now be discounted using an "upper-medium grade (low credit-risk)" yield. The accepted interpretation of this requirement is that insurers should use exclusively A-rated bonds when constructing the discount curve. The liability value changes driven by quarter-over-quarter updating of this curve will flow through accumulated other comprehensive income (AOCI), which is comparable to the existing treatment on available-for-sale fixed income investments.

This new requirement contrasts with existing guidance, which prescribes a rate based on the insurer's estimate (at inception of the contract) of the anticipated investment yield on the underlying asset portfolio over the life of the contract. Additionally, LDTI does not include a factor for potential adverse deviation to cash flow assumptions.

The goal of this research note is to examine in more detail the process of constructing discount curves in accordance with LDTI, including the criteria that should be used, and to assess the quality of the different algorithms used to build out a complete curve (including interpolation, extrapolation, and smoothing). We then assess the quality of several curve-fitting approaches by examining the fit to bond data for several historical dates covering a range of market conditions. Finally, we compare one of these curve construction methods with a more simplistic approach leveraging Treasury rates and average spreads, and examine the implications for the GAAP earnings and balance sheet for a cohort of payout annuities. First, though, some preliminaries.

## What are A-rated bonds?

As just stated, the industry consensus is that "upper-medium-grade" means A-rated credit quality. The three major bond credit-rating companies, S&P Global Bond, Moody's Investors Service, and Fitch Ratings, each have a category of A-rated bonds with similar language describing the quality of such bonds:

1. **S&P:** "Strong capacity to meet financial obligations, but somewhat susceptible to adverse economic conditions and changes in circumstances."
2. **Moody's:** "Obligations of upper-medium-grade, with low credit risk."
3. **Fitch:** "Strong capacity to meet financial obligations, but somewhat susceptible to adverse economic conditions and changes in circumstances."

Insurers will want to mark their liabilities for AOCI using discount rates that are consistent with the market valuation of the supporting assets, subject to the constraints of the LDTI discount rates definition. This would include reflecting the level and slope of risk-free rates, and the level and slope of spreads, among other possible dimensions that contribute to curve shape. Additionally, the discount curve will have to extend beyond the availability market data, which tends to be spotty after 10 to 15 years, whereas liability cash flows can extend for 30 or more years. The chart in Figure 1 shows the universe of available bonds at different maturities.

**FIGURE 1: CORPORATE A-RATED BOND MARKET YIELD VS. MATURITY AS OF JUNE 30, 2019**

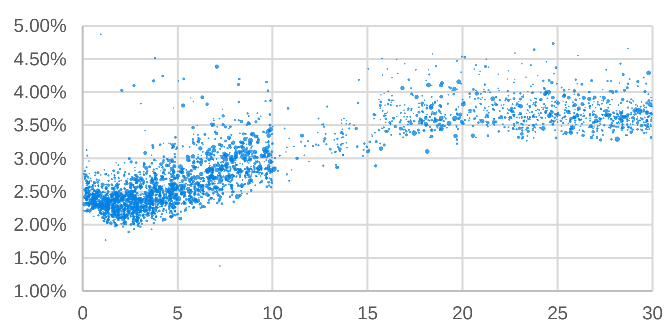

Any approach should use liquid instruments to avoid distortions caused by stale pricing. Also, it should reflect bond market characteristics, but generate a relatively stable curve that also captures market movements. The approach to extrapolation should also ensure that projections are consistent with yields before 30 years and with long-term investment returns available in the market. Finally, long-term rates beyond the last liquid point (30 years for the United States, though this is credit-dependent) should grade to a stable long-term forward rate assumption, which will promote a stable long-term basis for the associate discount curve.[1]

Discount curve construction should also conform to standard best practices, which include:

1. It should be transparent in construction, using accepted, logical, and accurate techniques and assumptions, enabling rates to be accepted by insurers, regulators, and auditors.
2. It should accurately reflect the high-grade corporate bond markets while minimizing idiosyncratic impacts of individual issues, and draw upon data that are as broad and as deep as possible at each maturity.
3. It should use readily and reliably available data that are updated frequently, and likely to remain available over an extended period.
4. The resulting spot yield curves should be smooth over the maturity range (smoothness often implies loss of fidelity to market prices, but it also improves explanatory power and attributions).
5. The curves should also evolve smoothly over time, so that they reflect changing conditions in established financial markets without inducing excessive short-term market volatility.

And the model used for building the curve must be well-behaved:

1. It should be parsimonious, that is, the model should use the minimum number of parameters to achieve the desired characteristics.
2. It should fit the data well.
3. It should be stable.
4. It should be flexible.
5. It should be consistent with expected behavior.
6. It should produce smooth, forward rates.

Using the above criteria, we have selected three different models to construct the discount curves: the Nelson-Siegal model, the Cairns model, and the Treasury High-Quality Market (HQM) methodology.

## Nelson-Siegel model

The Nelson-Siegel model[2] parameterizes the forward rates, $f(t)$, as follows: $f(t) = \beta_0 + \beta_1 \cdot e^{-\frac{t}{\tau}} + \beta_2 \cdot \frac{t}{\tau} e^{-\frac{t}{\tau}}$. This model has four parameters: $\{\beta_0, \beta_1, \beta_2, \tau\}$, three of which are usually free, and the fourth, $\tau$, specified. In our calibration, however, we constrained, but did not fix $\tau$. Figure 2 shows the three component factors of the Nelson-Siegel model. The forward curve is a linear combination of these factors, where the factor loadings may be positive or negative.

FIGURE 2: NELSON-SIEGEL COMPONENT FACTORS (T = 2.5)

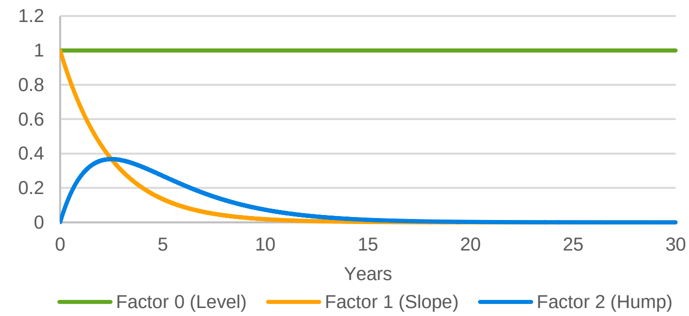


## Cairns model

The Cairns model[3] parameterizes the forward rates, $f(t)$, as follows: $f(t) = \beta_0 + \beta_1 \cdot e^{-c_1 \cdot t} + \ldots + \beta_n \cdot e^{-c_n \cdot t}$. Following Cairns (1998), we have used n = 4 and fixed $\{c_1, c_2, c_3, c_4\} = \{0.2, 0.4, 0.8, 1.6\}$. This leaves four free parameters: $\{\beta_1, \beta_2, \beta_3, \beta_4\}$. Figure 3 shows each of the component factors in the Cairns model.

FIGURE 3: CAIRNS MODEL COMPONENT FACTORS (N = 4, CI = [0.2, 0.4, 0.8, 1.6])

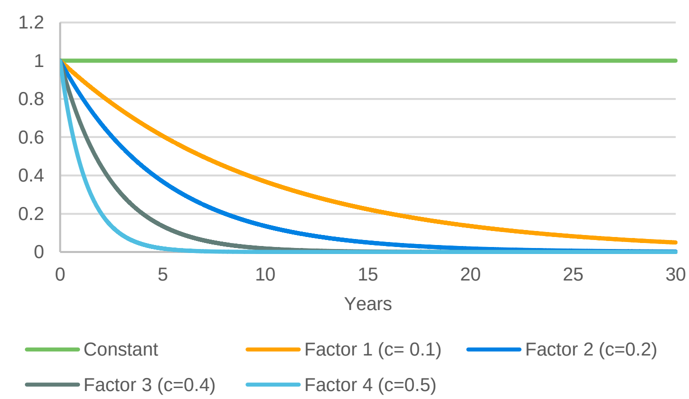


[1] This point might be controversial, but we believe that grading forward to a long-term assumption has several advantages. First, it reduces nonlocal impacts under some interpolation methodologies (such as cubic splines). Second, it promotes stability in valuation past the point for which there is market data, which reduces the extent to which the methodology is valued instead of the liability. Of course, there is a risk that inappropriate long-term assumptions are used to understate the true liability value, so a formalized process to setting these assumptions is required.

[2] Nelson, Charles R. & Siegel, Andrew F. (1987). Parsimonious Modeling of Yield Curves.

[3] Cairns, Andrew J. G. (1998). Descriptive Bond-Yield and Forward-Rate Models for the British Government Securities' Market.

## Treasury HQM methodology

The U.S. Treasury publishes yield curves[4] for the high-quality market (HQM) of corporate bonds (rated AAA, AA, or A), for purposes of discounting pension liabilities. The Treasury estimates such curves using clamped splines fitted to bond data. Specifically, forward rates are equal to a linear combination of five cubic splines, constructed to satisfy the following constraints:

1. Knots are fixed at 0, 1.5, 3, 7, 15, and 30 years.
2. First and second derivatives are continuous at the interior knots (1.5, 3, 7, and 15 years).
3. Boundary conditions are imposed at both ends of the 30-year fitting window.
4. The time zero forward rates are linear (second derivative equals zero).
5. The forward rates at 30 years are flat (first derivative equals zero).
6. The forward rate at 30 years equals the average of the 15-year and 30-year forward rates.

Figure 4 shows the basis splines resulting from these constraints.

**FIGURE 4: TREASURY HQM CONTRAINED B-SPLINES (KNOTS = [0, 1.5, 3, 7, 15, 30])**

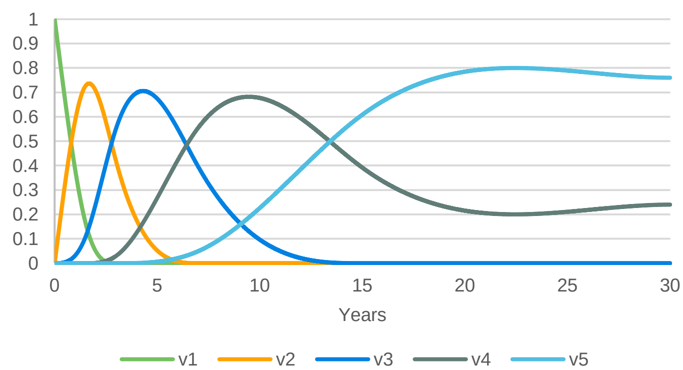


## Fitting procedure

Each of the parametric interest rate models described above was fit using nonlinear regression on individual corporate bond prices at a given date. The Huber loss function was used to minimize the impact of outliers.[5] To account for variation in bond size, the residuals were weighted by the square of par value. To compensate for heteroscedasticity[6] of prices, the residuals for bonds with durations greater than one year were further weighted by the inverse of the square of duration.

In order to fit on prices, one needs to estimate a model price (and a residual) for each bond and a given parameter estimate. This was done as follows:

1. For each bond in the data set, a vector of cash flows was estimated based on the coupon rate and maturity date. This vector has daily time steps with a 30/360-day count convention.
2. For a given estimate of the parameters of one of the models, a vector of forward rates is calculated.
3. The vector of forward rates is converted to discount factors.
4. The discount factors of step 3 are applied to the cash flows of step 1 to arrive at a model price
5. The (unweighted) residual is the difference between the model price of step 4 and the actual (invoice) price.

## Data

For this, we used ICE's fixed income index data service. The data from ICE has a number of fields, including par amount, coupon rate, maturity date, yield-to-maturity, effective yield, clean price, and accrued interest. Unfortunately, the data set from ICE did not include a payment frequency, so we assumed semiannual coupons, which is standard.

The data set from ICE includes all U.S. dollar-denominated corporate bonds, with issuances above $250 million. For this exercise, bonds were filtered according to the following criteria:

1. Only bonds with ratings of A1, A2, or A3, as assigned by ICE based on nationally recognized statistical ratings organization (NRSRO) ratings, were included.
2. The raw data was already filtered to include only bonds with outstanding par greater than $250 million.
3. Bonds with a maturity of greater than 30 years were excluded.
4. Bonds with a maturity of less than three months were excluded.
5. The data set does not include any fields to indicate floating rate bonds or callable bonds. To remove such bonds, bonds with yield-to-maturity not equal to effective yield were discarded.

[4] Girola, James A. (2007). The Corporate Bond Yield Curve for the Pension Protection Act.

[5] Note that, when using any type of robust loss function, such as the Huber function, loss does improve stability and reduce squiggliness in calm markets. However, it also biases the curve a bit lower in stressed markets because bonds with blown-out spreads are treated as if they are outliers and make reduced contributions to the fitted curve. For high-quality bonds, such as A-rated bonds, it does not make a huge difference, but for late 2008 this was worth about 200 bps in the short end of the BBB curve.

[6] Heteroscedasticity is when the standard errors of a variable are nonconstant over time.

# Sample results

The rest of this memo presents the results for a selection of dates intended to cover a variety of yield environments. The dates are summarized in the table in Figure 5.

**FIGURE 5: DATES AND YIELD ENVIRONMENTS**

| Month End | Yield Environment | 2Y UST | 10Y UST | 30Y UST | Avg Corp A Yield | Avg Corp A OAS | Std Dev Corp A OAS |
|---|---|---|---|---|---|---|---|
| Nov 2006 | Inverted Treasuries | 4.62% | 4.46% | 4.56% | 5.43% | 0.79% | 0.31% |
| Dec 2008 | Extremely stressed | 0.76 | 2.25 | 2.69 | 7.51 | 5.62 | 3.36 |
| Dec 2013 | Steep yield curve | 0.38 | 3.04 | 3.96 | 2.99 | 0.93 | 0.50 |

**FIGURE 6: FITTED RATES AND YIELDS, 2006-2013**

**NOVEMBER 2006**

**FITTED PAR RATES VS. MARKET YIELDS**

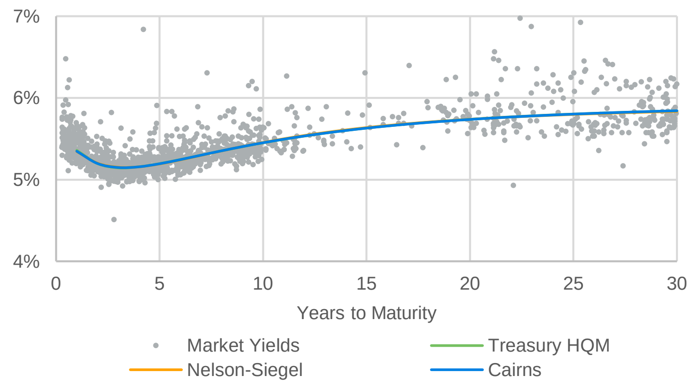


**FITTED FORWARD RATES**

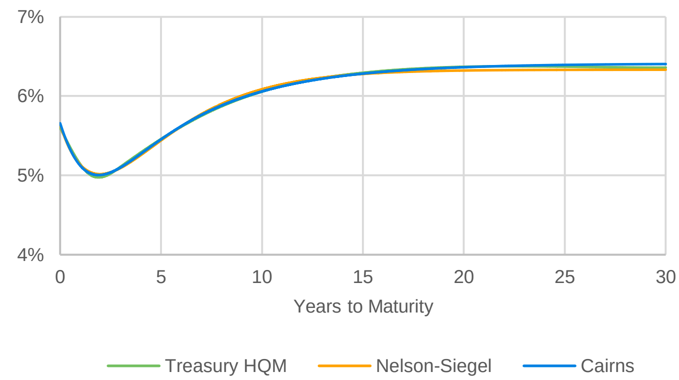


**FITTED ZERO COUPON YIELDS**

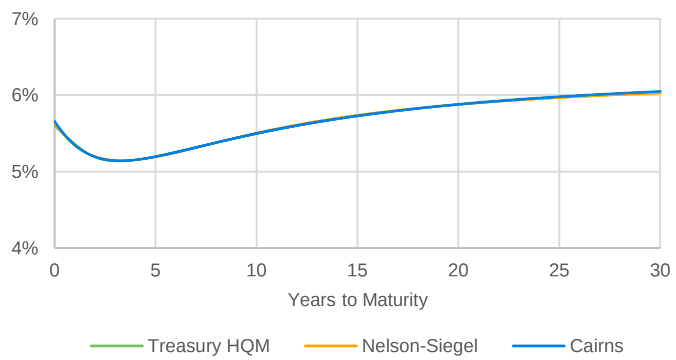

## DECEMBER 2008

### FITTED PAR RATES VS. MARKET YIELDS

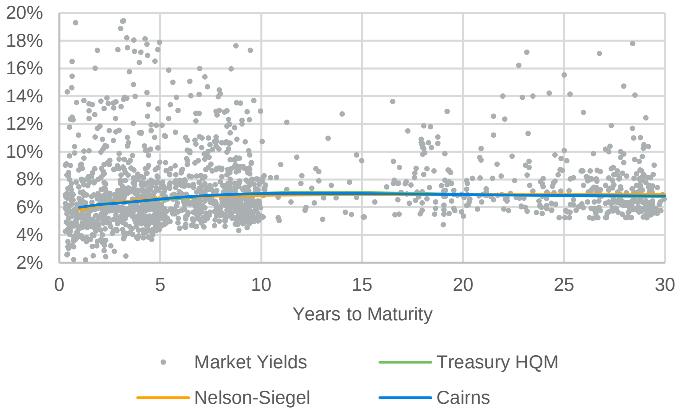


### FITTED ZERO COUPON YIELDS

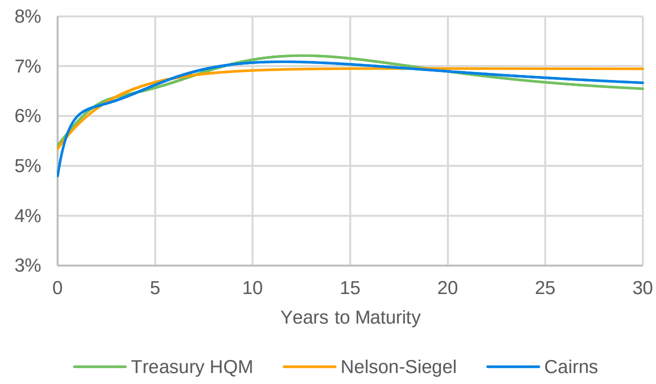


### FITTED FORWARD RATES

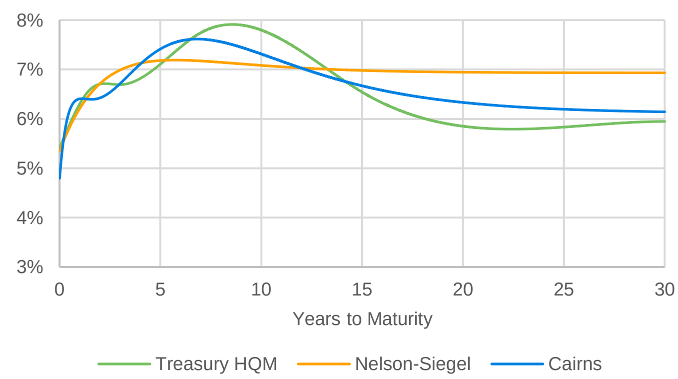


## DECEMBER 2013

### FITTED PAR RATES VS. MARKET YIELDS

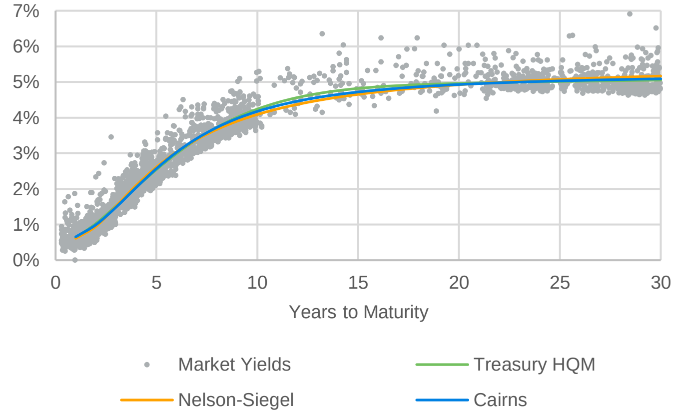


### FITTED ZERO COUPON YIELDS

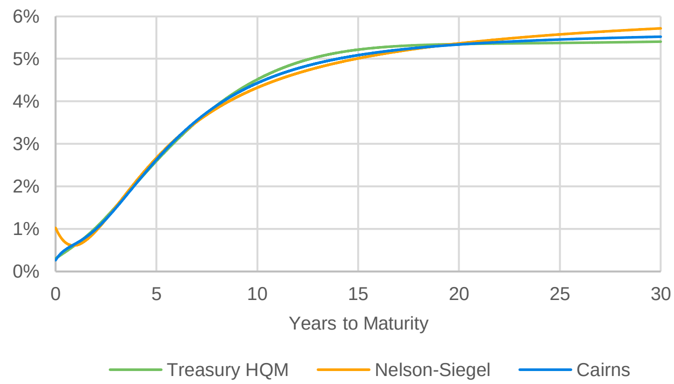


### FITTED FORWARD RATES

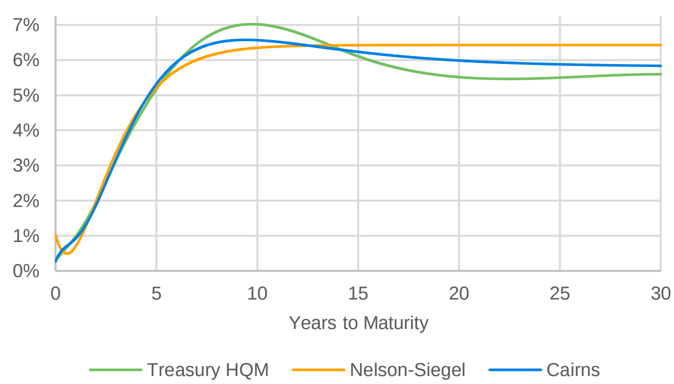

### OBSERVATIONS ON DIFFERENCES BETWEEN PARAMETRIC FORMS

1. All four parameterizations tend to produce similar par curves, which appear to be a good fit to the central tendency of the market yields in the fitting data sets. The greatest dispersion in the fitted par yields usually occurs where the data is thinnest, between 10 and 15 years.
2. All four parameterizations tend to produce similar zero curves, but there can be small differences especially at the short end and the long end of the curve. For the dates shown the difference in the forward rates between any two parameterizations was 10 to 40 basis points (bps).
3. The Treasury HQM methodology tends to produce forward rate curves with a hump in the 10-year to 15-year range, which is not observed in the other methodologies. This is likely a side effect of the boundary conditions and the fixed knot points.
4. The Cairns parameterization is prone to wiggly short-term forward rates (one to three years).
5. The Nelson-Siegel parameterizations sometimes produce a bit of a dip in the short-term forward rates. These dips may reflect actual market data, as they show small downward trajectories, though there is too much noise in the data to conclude that is an actual feature being captured by the model, as opposed to a feature of the algorithm.

## Implications for GAAP earnings and balance sheet

Next, we examined the impact of different curve approaches on GAAP earnings and balance sheet for a single cohort of annuity business. We compared the curve-fitting approach to a simplified Treasury forward plus Option-Adjusted Spread (OAS) approach. For the curve-fitting approach we have illustrated the Treasury HQM formulation. However, in light of the results noted above, the specific formulation does not have a large impact on the results.

We applied each of these curves to a portfolio of single-premium immediate annuities with a weighted average life (WAL) of 12 years, which is approximately representative of a pension-risk transfer case or a cohort of individual retirement annuities.

We used actual bond data and Treasury yields over the period from June 1997 to June 2017 to simulate the results. Figure 7 shows the pattern of cash flows and the forward rates for each approach as of June 1997.

FIGURE 7: COMPARE FORWARD RATES

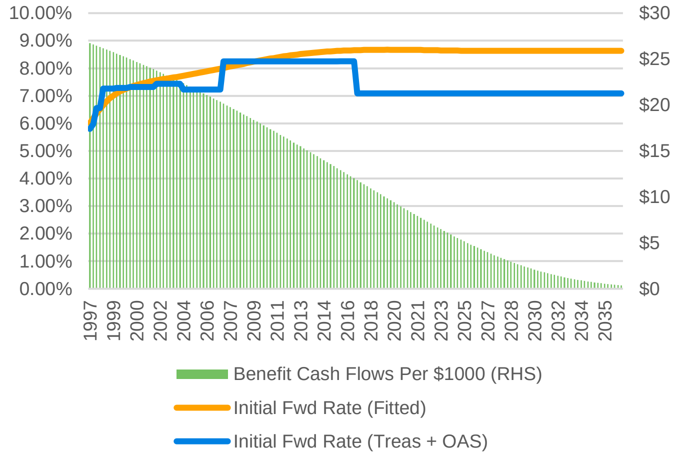


As you can see, the term structure under the Treasury plus OAS is substantially flatter than the curve built off A-rated bonds. This is not surprising, given that the former approach would tend to understate longer-duration credit spreads.

We evaluated the profitability pattern under two alternate assumptions for locking in the at-issue discount rates: first, locking in the forward curve, and second, locking in an equivalent single rate that produces the same present value of benefits at issue.

### RESULTS UNDER LOCKED-IN FORWARD CURVES

The chart in Figure 8 illustrates the GAAP return on assets (ROA)—earnings excluding other comprehensive income (OCI) divided by the book value of invested assets—for each discounting approach, under the assumption of locked-in forward curves:

FIGURE 8: GAAP RETURN ON ASSETS

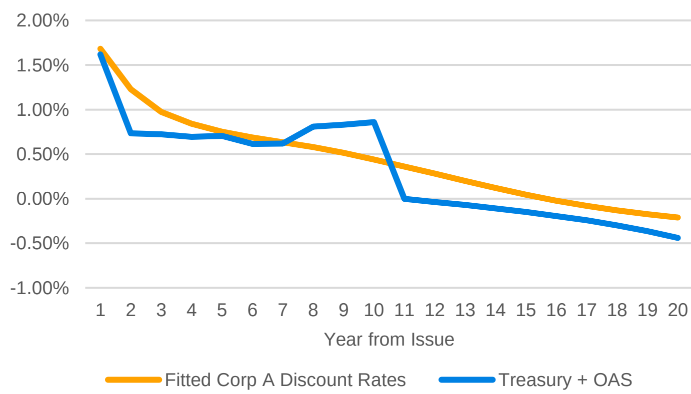

For either approach, locking in the forward curve will front-load profitability when the forward curve is upward-sloping, as is usually the case. Comparing the two approaches, the profitability pattern is relatively smooth for the fitted approach, and non-smooth for the Treasury plus OAS method, which is due to interpolation differences.

Figure 9 illustrates how these curves translate into dollars of GAAP operating income.

FIGURE 9: GAAP OPERATING INCOME PER $1,000 OF PREMIUMS

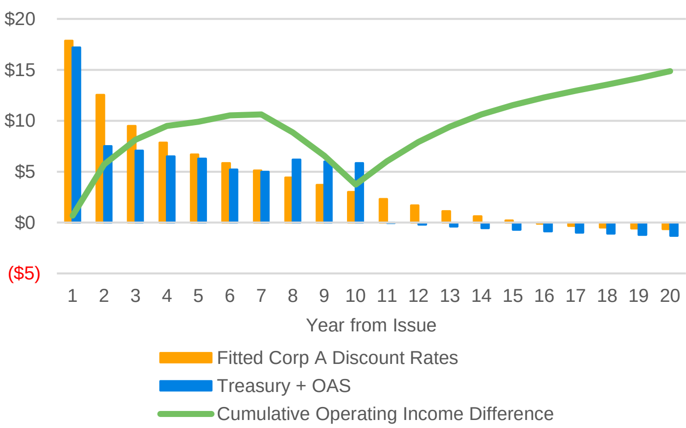


A single point in time is useful for understanding the first-order impacts of discount curve choices, but how the different methodologies behave over time is just as important. Figure 10 shows the mark-to-market (MTM) impact assuming a backing portfolio that is to be purchased at par value with a runout pattern exactly matching the expected runout of statutory reserves and capital (such a portfolio is approximately cash flow and duration matched). The portfolio is assumed to be invested in 50% A-rated and 50% BBB-rated corporate bonds, with coupon rates equal to average market yields at issuance of the annuities.

FIGURE 10: GAAP AOCI (DUE TO NET ASSET/LIABILITY MTM) PER $1,000 OF PREMIUM

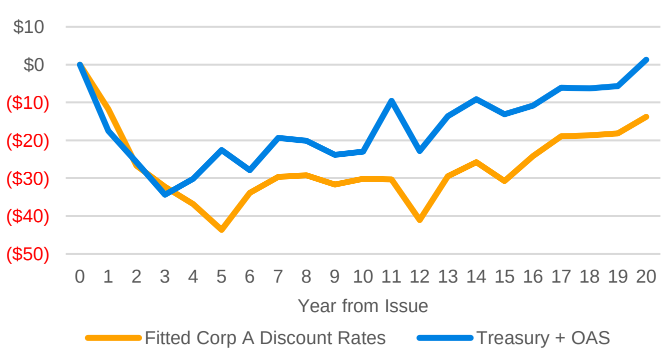


Under the locked-in forward curve approach, there is a bias toward negative AOCI. As the business rolls down the initial forward curve the average discount rate is increasing. If the actual rate environment is static, that will lead to liability "book values" that are lower than liability "market values," which contributes to negative AOCI. Of course, the rate environment will not be static, but this effect will exist as long as bond yields at issue are upward sloping and the upward slope persists to some extent after issue.

### RESULTS UNDER LOCKED-IN EQUIVALENT SINGLE RATES

The chart in Figure 11 illustrates the GAAP ROA under the assumption of locked-in single equivalent rates.

FIGURE 11: GAAP RETURN ON ASSETS

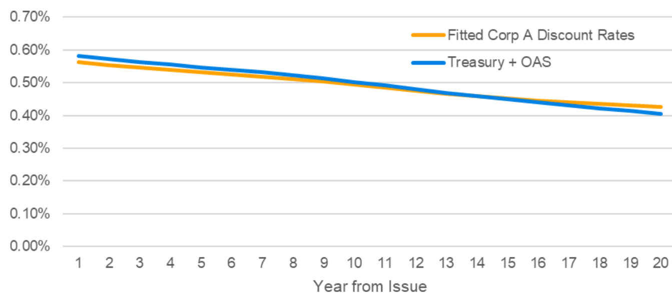


When a single equivalent rate is locked in for calculating earnings, instead of the full forward curve, the differences between the two approaches become significantly less pronounced. This is due to several factors:

1. The simple Treasury plus OAS approach captures much of the same information about the absolute level of interest rates that is captured by the curve-fitting approach. The notable exception is the additional spread commanded for longer-term bonds, which are less represented in the average OAS. However the cash flows used in this example are not long enough for this to make a significant impact on the results.
2. Using a single equivalent rate eliminates differences in the shapes of the curves produced by the two approaches.
3. Under the GAAP accounting framework, differences in the benefit reserve levels are offset by deferred profit liability (DPL), which will tend to equalize the earnings patterns for small differences in discount rates (lower initial reserves lead to higher initial DPL; the higher reserve accretion rate is offset by greater release of DPL).

We also illustrate this more subdued impact on operating income in Figure 12.

**FIGURE 12: GAAP OPERATING INCOME PER $1,000 OF PREMIUMS**

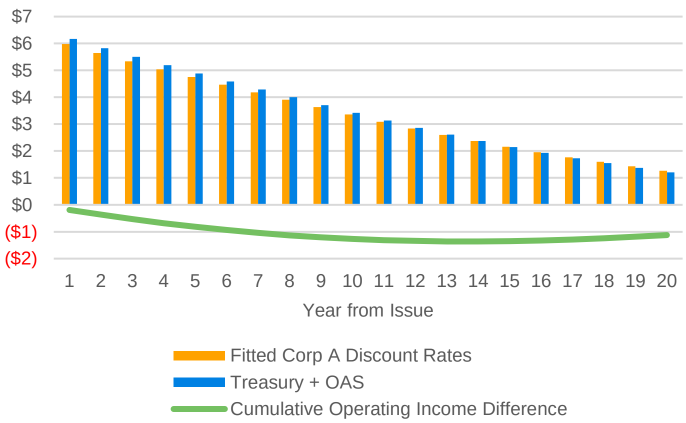


Figure 13 shows the net AOCI when a flat rate is locked in for bifurcating earnings and OCI.

**FIGURE 13: GAAP AOCI (DUE TO NET ASSET/LIABILITY MTM) PER $1,000 OF PREMIUM**

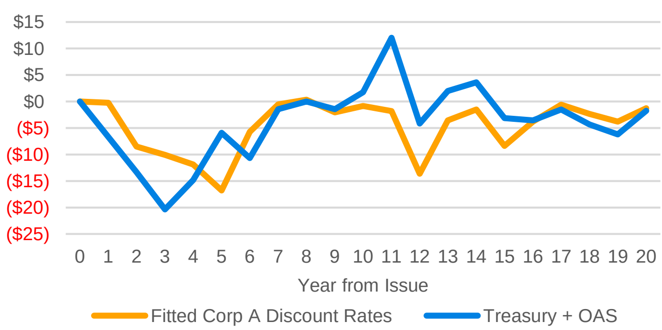


The bias toward negative net AOCI does not exist when a single rate is used for the lock-in.

### CONCLUSIONS FROM ILLUSTRATION

The decision as to whether to lock in a curve or an equivalent flat rate makes a significant difference in the pattern of profitability. When a curve is locked in, the shape of it and thus the interpolation and smoothing methodology drives the profitability pattern. When an equivalent flat rate is used, the interpolation and smoothing methodology makes little difference.

### FUTURE INVESTIGATIONS

In this paper, we have investigated how different curve parameterizations affect the valuation of nonparticipating products under GAAP LDTI. If this were all that was required, then implementation would be straightforward. However, a few key items will also be critical when setting up the discount curve process. Further contributions to this topic could address the following:

1. How much should insurers use internal asset data instead of external vendor data, when constructing discount curves? Internal data creates better asset-liability management (ALM) alignment, but potentially sacrifices full coverage of the bond market. If an insurer takes a blended approach, it should be explicit about where and why it is augmenting its assets with external bond information.
2. To what extent should insurers include a mix of credit quality in curve construction? As mentioned above, LDTI does not require the use of only A-rated bonds, so an insurer can substitute a blended portfolio.[7] This may better align with the insurer's actual investments, but may also create more volatility, and induce greater model-dependent behaviors.
3. Should insurers create liability-specific curves? What are the benefits/costs of such an approach?

GAAP LDTI will create a sea change in how insurers are valuing their liabilities. While we believe this will benefit all stakeholders in the long run, the short-term challenges are not trivial. In this paper, we explored one of these challenges: how to construct a discount curve that reflects an upper-middle credit quality. We looked at what considerations insurers should use when deciding on a parametrization, the consequences of these choices for common approaches, and the resultant liability valuation impact, both on a single date and over multiple reporting periods. We hope this will help insurers think through how they will manage this part of their implementation journey.

### CONTACT

Josh Dobiac
josh.dobiac@milliman.com

Bill Matczak
bill.matczak@milliman.com

Jeff Greco
Jeff.greco@milliman.com

[7] However, it should be mentioned that the flexibility here is contingent upon auditor approval, and we have heard anecdotally that some auditors are being fairly strict with requiring the inclusion of only A-rated bonds.