Yield Curve Extrapolation Methods

Methodologies for Valuing Liability Cash Flows That Extend Beyond the Maximum Yield Curve

March 2019

SOCIETY OF ACTUARIES

Yield Curve Extrapolation Methods

Methodologies for Valuing Cash Flows That Extend Beyond the Maximum Yield Curve

AUTHORS

Kemi Akinyemi, FSA, EA, MAAA
Actuarial Consultant
Risk & Regulatory Consulting

Jack Kerbeshian, FSA, MAAA
Actuarial Consultant
Risk & Regulatory Consulting

Benjamin Leiser, FSA, MAAA
Consulting Actuary
Risk & Regulatory Consulting

Patricia Matson, FSA, MAAA
Partner
Risk & Regulatory Consulting

SPONSOR

Society of Actuaries Committee on Finance Research

CONTENTS

Yield Curve Extrapolation Methods

Methodologies for Valuing Cash Flows That Extend Beyond the Maximum Yield Curve

Executive Summary

The Society of Actuaries Committee on Finance Research engaged Risk & Regulatory Consulting LLC (RRC) to conduct research on methodologies used for yield curve extrapolation to value liability cash flows that extend beyond the maximum observable portion of the yield curve.

The researchers performed research on the methods available in theory and used in practice and summarized that information within this report. They also developed and provided questionnaires to various experts within the industry to comment on methods used to extrapolate the yield curve both in practice and in theory, and they consolidated the results of the independent research and the findings obtained from the verbal and written responses to the questionnaires to provide a comprehensive view of the yield curve extrapolation methods.

There are two philosophical approaches to yield curve extrapolation: one that emphasizes market consistency at a point in time, and another that emphasizes liability stability across time. The advantages and disadvantages of each approach are explored in this report.

The initial goal when extrapolating the yield curve under many methods is to determine an ultimate long-term forward rate (Definitions) (UFR) to which the observable yield curve will converge.

Some of the simpler extrapolation models include the Simple Monopole or Dipole methods, the Flat Rate extrapolation method and the Linear First-Order (Definitions) extrapolation method, along with a few other first-order extrapolation methods. The more complex models include the Nelson-Siegel/Svennson, Smith-Wilson, Cairns and Cubic Spline methods. The benefits and drawbacks of these methods are covered in the report.

This report compares the Nelson-Siegel, Svennson and Smith-Wilson methods in particular and provides expert panelist opinions on extrapolation methods. This includes industry approaches, methods they use, determining the UFR, duration of UFR, speed of convergence to the UFR, shape and smoothness of transition from observed rates to the extrapolated rate, mechanics or processes used to fit the curve, complexity of models and selection of software.

Based on the research, it appears that most of the methods that are analyzed in theoretical discussion are similar to what practitioners use. Although the assumptions that feed into an extrapolation method often have a greater impact than the technical methodology itself, the choice of the method does have an impact on the results. For example, using the current year forward rate extrapolated out into the future will have a much different result than a method that grades over time.

Many of the panelists stated that their method is simple and adequate. They believe that other methods involve more complex math without much evidence that they are actually any more theoretically justifiable. A simple approach might be used because the liabilities involve nonguaranteed elements, which minimizes the effect of any choice of approach as long as minimum guarantees are not in the money. Others prefer a

complex method, perhaps using a Cubic Spline Nelson-Siegel method that is at the more sophisticated end of the spectrum. Although they may have investigated many other techniques, they could believe this is the best solution given their requirements for highly accurate fits and an automated production process. One panelist stated they endeavor to use the simplest model possible, but no simpler than what is necessary to be consistent with the market and economic principles. Simplicity is not necessarily dependent on the fitting of the model.

Section 1: Introduction

The Society of Actuaries (SOA) Committee on Finance Research sponsored this research study (hereafter "the Study") to investigate yield curve extrapolation methods used to value cash flows that extend beyond the last liquid point of a yield curve. The researchers carried out the main objectives of this project, including background research, interviewing experts and developing this report. The objectives of the Study were to do the following:

1. Perform research regarding current industry methods and approaches for extrapolation of the yield curve beyond the current investible universe, including details of how they are applied, and include any observations we have on the benefits and drawbacks and prevalence of their use. The primary focus of the research is on the U.S. and Canada; however, other international methodologies and guidance around extrapolating the yield curve will also be reviewed to the extent they are directly applicable to the U.S. and Canada.
2. Identify and interview a broad group of subject matter experts with strong industry representation to avoid any bias in the results.
3. Supplement the research with results of the interviews, including any theoretical and practical issues noted with methods. Responses will likely be too varied to draw strong conclusions from the interviews alone because yield curve extrapolation methods are not at the point of a significant convergence of practice.
4. Summarize the results of the research and interviews, including the approach, the information gathered and the conclusions generated. The summary includes the following:

    - A description of the methods used in the U.S. and Canadian insurance industries
    - Additional commentary on methods used in certain significant regulatory regimes
    - Commentary on the applicability of the methods beyond the insurance industry, specifically with respect to valuation of pension and other postretirement benefits
    - Benefits and drawbacks for each of the approaches, both practical and theoretical
    - Other considerations for U.S. and Canadian actuaries trying to implement an approach for valuing insurance and pension liabilities

Based on the results of the research, we have summarized the approach, information gathered and conclusions. The summary includes information that is responsive to each of the objectives outlined above.

Risk & Regulatory Consulting (RRC) conducted this research. Information regarding RRC can be found on their website, www.riskreg.com .

One of the most fundamental concepts in actuarial practice is the time value of money. For any work in which future cash flows are allowed for, such as reserving or pricing, it is natural to discount to present values so that an appropriate amount of money can be set aside today, allowing for future investment returns.

Risk-free yield curves are the basic building blocks for the valuation of future financial claims and long-term risk management work. Despite their fundamental importance, it turns out that measuring and estimating suitable risk-free interest rates present major challenges. In highly developed fixed-income markets, we may be able to observe bonds or interest rate swap contracts with maturities of up to 50 years. In less developed markets, liquid bond quotations might be limited to only a few years. In exceptional circumstances, there may be no traded risk-free instruments at all. Of course, the liabilities of long-term financial institutions frequently extend beyond the term of available market instruments. To value these long-term claims and assess risk, practitioners must extrapolate yield curves to generate a set of "prices" for the assumed, inferred prices of discount bonds beyond the term of the longest available traded cash flow. A good yield curve estimation method must deliver extrapolated curves that are credible at a single point in time and where changes over time in extrapolated rates can be justified.

Yield curve construction work requires completing two fundamental tasks: first, collating market data and fitting a continuous curve to the term of the longest available market instrument, and, second, extrapolating from the longest available market data toward some long-term assumption for forward interest rates.

Extrapolation also requires answering two questions about the path of forward interest rates beyond the longest market data point:

- First, what is an appropriate assumption for the infinite-maturity, unconditional forward rate of interest?
- Second, what path is chosen between the longest (smoothed) market forward rate and this long-term rate? In particular, the analyst needs to determine the speed at which the extrapolated forward rate tends toward the long-term asymptote.

The documentation below outlines a summary of the yield curve extrapolation methods, with additional commentary provided by the subject matter experts. We performed a search of commonly used yield curve extrapolation methods based on various research papers. The concept of yield curve extrapolation was similar across various documents, and we give a summary of this analysis.

In extrapolating the yield curve beyond the current investable universe, several considerations are important, including market consistency at a point in time, liability stability across time and the smoothness of the extrapolated curve.

Section 2: Methodology

2.1 Initial Research

The researchers conducted an initial review of the existing literature regarding industry approaches for extrapolation of the yield curve, considering the purpose of the measurement, which is to value liability cash flows that extend beyond the maximum yield curve. The researchers reviewed several research documents covering current approaches used in both the U.S. and Canada, as well as international approaches as applicable. The papers and research documents used for this purpose are included in the References of this report.

2.2 Questionnaire and Interviews

The researchers developed an interview questionnaire based on conducting an initial review of the literature and on additional questions to obtain related information from subject matter experts. The questionnaire covered topics such as the following:

- Discussion of industry approaches for extrapolating the yield curve and the situations (specific products, specific applications) in which each is used
- Key assumptions and mechanics considered in the extrapolation of the yield curve
- Benefits and drawbacks of the various approaches
- Practical challenges that arise from various methods
- Regulatory approaches for extrapolation, including the rationale for the choice(s)
- Governance and standards of practice

Twelve subject matter experts were interviewed, and, although several experts opted for a written response, live interviews were also conducted to discuss the questionnaire or commentaries to the questionnaire.

The subject matter experts who were interviewed had one or more of the following characteristics:

- Practical experience in extrapolating the yield curve for purposes of valuing long-term liabilities (for example, for purposes of internal capital calculations) across a range of life, annuity, health, property and casualty and pension liabilities
- Experience in the development and use of economic scenario generators used to value long-term insurance liabilities
- Background in quantitative finance
- Experience researching and developing theoretical methodologies for extrapolation of the yield curve
- Experience developing and/or implementing regulatory requirements that involve yield curve extrapolation (such as Solvency II)
- Experience in a range of jurisdictions, including the U.S., Canada and Australia
- Background that encompasses a broad number and type of companies, to get a broad and representative understanding of approaches used in the U.S. and Canada today.

The panelists provided their opinions relating to the extrapolation of the yield curve beyond the investable universe by responding to multiple questions within the questionnaire in writing, verbally or both. The researchers then consolidated all questionnaire responses. Since the sample size was small, and the questions were open-ended, the responses were not conducive to present a distribution of results.

2.3 Panelists' Backgrounds

The panelists' backgrounds included both theoretical and practical experience, with more than half the panelists having both. Many have used a range of yield curve construction and extrapolation techniques over many years. These panelists have also published many papers on the topic.

All of our panelists have experience, including the development and use of economic scenario generators (ESGs). This includes experience building generators and using various ESGs and yield curves in valuation, pricing and own risk solvency assessment work. The panelists have also often supervised the use of numerous generators, including both proprietary and industry standard. One panelist manages a market-leading ESG with a large worldwide customer base. This customer base spans a range of industries, liability types and regulatory regimes. Yield curve construction is a key input into the ESG calibration for all their ESG clients.

The range of liability categories applied using the extrapolation of the yield curve is wide. This includes life insurance, annuities, long-term care, guaranteed investment contracts, group annuities, property and casualty (which includes some long-term portfolios such as workers' compensation, lifetime care and mortgage insurance), pension and health.

The panelists' experience includes developing and/or implementing regulatory requirements that involve yield curve extrapolation. Some have been members of committees that have set professional standards for model calibration. Other experience includes the following:

- Working with a participating company in the International Association of Insurance Supervisors (IAIS) field test of the proposed international capital standard
- Producing yield curves on multiple bases, including prescribed methodologies such as Solvency II and the Swiss Solvency Test
- Producing papers on yield curve construction and actively participating in the insurance industry's lobbying efforts in the Solvency II consultations
- Producing balance sheet reserve estimates often requiring market-consistent risk-free rates, which involves a judgment on risk curve extrapolation
- Implementing requirements, namely, supervising the modification and use of rate models based on the American Academy of Actuaries' generator and VM-20 (Requirements for Principle-Based Reserves for Life Products) requirements.

2.4 Summary of Research and Results of Questionnaire

The researchers summarized the research and enhanced the results of the research based on the subject matter experts' written and verbal responses to the questionnaires.

Section 3: Summary of Research Results

3.1 Philosophical Considerations: Market Consistency/Liability Stability

Broadly speaking, there are two philosophical approaches to yield curve extrapolation: one that emphasizes market consistency at a point in time, and another that emphasizes liability stability across time.

With the market consistency approach, the aim of yield curve extrapolation is to estimate what the market price of longer-term assets would have been on a particular day if those assets actually existed and were traded in the financial markets. The emphasis here is on producing a liability value that is equivalent to the value that would be required to transfer that liability in current market conditions.

On the other hand, those adopting an approach emphasizing liability stability believe that the volatility introduced by requiring market consistent prices at a point in time is unhelpful for running a business over many years. They also believe that since insurance liabilities are rarely traded, the idea of identifying market value is less relevant. Some proponents of this view may be supportive of the idea of valuing liabilities for which hedging instruments are available with reference to the prices of those instruments. However, where those instruments do not exist, or are not sufficiently liquid, they would rather emphasize liability stability over market consistency.

Each approach has its advantages and disadvantages. Although an approach that emphasizes liability stability may be superficially attractive, some believe it may cause insurers to underestimate the true economic cost of writing long-term insurance contracts. It would not be in the firm's, policyholders' or shareholders' future interests if long-term contracts were issued too cheaply. In addition, an artificially stabilized liability valuation may also make it harder to hedge balance sheet risk over the longer term.

On the other hand, it has been suggested as part of the Solvency II debate that emphasizing market price consistency may play a role in magnifying economic or financial downturns, or, in other words, it may add to pro-cyclicality. The Solvency II liability evaluation is a fair value approach, which seeks to capture changes in asset and liability values over time. It is an amount at which an asset could be exchanged between knowledgeable and willing parties in an arm's-length transaction.

Table 1 breaks down some differences between market-consistent (risk-neutral) and liability-stability (real-world) scenarios.

Table 1

Differences between Risk-Neutral and Real-World Scenarios

	Risk-Neutral Scenarios	Real-World Scenarios
View	Forward looking	Backward looking
Presumption	History is irrelevant	History will repeat
Asset return	Risk-free rate	Risk rate plus risk premium
Method	Objective mechanical	Subjective explicit judgment
Philosophy	Hard to beat market	Can beat market
Data	Snapshot of current market	Historical market data

3.2 Extrapolation and the Ultimate Long-Term Forward Rate

The Initial goal when extrapolating the yield curve under many methods is to determine an ultimate long-term forward rate (UFR) to which the observable yield curve will converge. The components of the UFR are the following (a+b+c-d):

a. Expected future inflation
b. Expected real short-term rate, which is the expected nominal short-term rate minus the expected future inflation
c. Term premia are the additional return an investor may expect as compensation for the longer-term investment and is the difference between the forward rate and the expected future short-term interest rate. The term premium acts as compensation for holding long-term bonds, whose value will fluctuate in the face of interest rate uncertainty, exposing the holder to mark to market losses. Term premia have the following components:

    i. Risk premia—investors demand a premium for locking in long-term investments. This acts as compensation for holding long-term bonds, whose value will fluctuate in the face of interest rate uncertainty, exposing the holder to mark to market losses.
    ii. Term preference—demand for long-term government securities from large institutional investors can drive down long-term forward rates because the long-term bonds offer a closer match to liabilities and are less risky investments to these investors.

d. Convexity (Definitions) effects—fixed-income investments have positive convexity, which can cause longer-term bonds to trade at higher values (lower yields). Convexity adjustment arises because of the nonlinear (convex) relationship between interest rates and bond prices.

Most of the change of the one-year forward rate is explained by change in the expected short rate. By year 10, most of the change is explained by change in the term premia. This is, in part, because short-term interest rates are expected to mean revert over time. In other words, expectations of future short-term interest rates are relatively fixed by durations of 10 years or more. On the other hand, the term premia appear to display little mean reversion.

Solvency II assumes that for most currencies the UFR will be 4.2%, a result of a 2% inflation rate and 2.2% long-term growth rate. Solvency II prescribes that the risk-free term structure, which is used to discount the best estimate liabilities, shall be derived from financial instruments for which a deep, liquid and transparent market exists. For maturities where such markets do not exist, the forward rates corresponding to the financial instrument should be extrapolated to the UFR to obtain the risk-free term structure. Each year at the end of March, the European Insurance and Occupational Pensions Authority (EIOPA) will calculate an effective UFR, and it will be used to extrapolate rate curves from January of the following year.

The pragmatic approach is to set the long-term term premia assumption to some assumed stable level. Although this may not be correct, it steers away from introducing potentially unwarranted volatility onto an insurer's balance sheet.

Some general principles are commonly used for estimation of the ultimate nominal forward rate. The rate should have stability where the unconditional nominal forward rate cannot be materially affected by short-term economic changes. There should also be consistency. If the long-term forward rate is approached at some very long-term horizon, this value is expected to be broadly the same around the world, and simplicity should be used. Although it is inherently difficult to estimate the long-term forward interest rates, the aim is to use a simple approach that is easy to understand that can be applied to a large number of economies.

The long-term yield curve should be built extrapolating from the last observable rate to the ultimate rate. Transition from observed rates to extrapolated rates should be smooth and exhibit continuity; however, this constraint can cause challenges with extrapolation. The variability or volatility of long-term forward rates should be lower than short-term rates (mainly due to mean reversion).

A couple of considerations for extrapolation include specifying the UFR and specifying the speed of convergence (not for all methodologies). Note that the speed of returning to the UFR varies by country, and the UFR is presumed to be reached when the slopes forward interest rate curve becomes zero.

The group of panelists have investigated or used both theoretically and practically the methods described below. When it comes to extrapolation, most consider only methods that effectively converge or revert to a specific instantaneous forward rate and that do not create arbitrage opportunities.

3.3 Sample of Extrapolation Methods

After determining the UFR, the next step is to determine the appropriate methodology for extending or extrapolating the yield curve beyond the current investable universe. Below we list several methodologies along with some detail on each method.

Industry approaches for extrapolating the yield curve, and the specific situations in which these approaches are used, include for valuations of both life insurance and annuities. Extension of the yield curve has been necessary when there was a requirement to report a "market" or "current" value for liabilities. Such values are used in GAAP accounting and can also appear in regulatory reports related to cash flow testing of reserve adequacy. For such purposes, it has been common to extend the yield curve very simply by freezing the last observable rate. The frozen rate may be the forward, spot or bond rate, and the other two rates are adjusted for consistency with the frozen or fixed rate.

One panelist stated extrapolation can be used by the product group for initially setting long pricing rates, and the approach is almost always very simple, such as freezing the forward rate or introducing a conservative (negative) bias to the forward rate. This panelist often works with spot rates (Definitions) because forward rates (Definitions) are hard to evaluate confidently (forwards are the derivatives of the spot curve and are sensitive to the smoothness of the spot curve).

The following paragraphs summarize the various methods researched.

The Simple Extrapolation Method is simple to implement. It has two variations:

- *The Simple Monopole method* assumes a constant single forward rate for all durations greater than 30.
- *The Simple Dipole method* uses the maximum observable (often 30-year) forward rate beyond that point.

See the paper "A risk management tool for long liabilities: The static control method" in the References for more information on these simple extrapolation methods.

The Flat Rate Extrapolation Method is similar to the simple extrapolation method. It assumes that the longest observable spot rate is extended infinitely throughout the nonobservable portion of the yield curve. For any extrapolation, the long rate is guaranteed to exist and to be finite; however, it will not remain constant across periods. The usage of the observable yield curve is small as the extrapolation relies entirely on the longest observable rate. The single factor driving the model is the longest observable rate. Therefore, the driving factor is a tradable quantity, although this could be limited when liquidity is low.

The Linear First-Order Extrapolation Method assumes that a first-order linear relationship exists between forward rates beyond the longest observable spot rate. The two factors driving the model are the gradient and scale factors. Where the two factors are determined exclusively by the observable yield curve, then they will be hedgeable. There is little risk of multicollinearity (Definitions) in this two-factor model, where one factor is influenced by the level of rates and the other is influenced by the slope. This method assumes that the forward rates beyond *M* years follow a first-order linear progression of the form

$$f_t(\tau) = a + b \times \tau, \tau > M,$$

where

- *a* and *b* are the parameters of the extrapolation, estimated via least squares (Definitions) and
- *M* represents the longest observable (and tradable) spot rate.

Other First-Order Extrapolation Methods

a. *The Power Spot Rate Extrapolation* assumes that forward rates beyond the longest observable spot rate follow a power relation (Definitions) This method assumes that the forward rates beyond *M* years follow a power progression of the form

    i. $f_t(\tau) = a \times \tau^b, \tau > M.$

b. *The Exponential Spot Rate Extrapolation* assumes that forward rates beyond the longest observable spot rate follow an exponential relation. This method assumes that the forward rates beyond *M* years follow an exponential progression of the form

$$f_t(\tau) = a \times e^{bxt}, \tau > M.$$

The long rate for these approaches either is zero or does not exist. The usage of the observable yield curve covers the latter end of the observable curve. The approach does not display intertemporal consistency (Definitions), because the long rate does not always exist. The two factors driving the model are the gradient and scale factors. Where the two factors are determined exclusively by the observable yield curve, they will be hedgeable.

The choice of nonlinear paths is partly a tension between simplicity and aesthetics, given that identifying the true shape involves considerable uncertainty. The linear path is reasonable, but more sophisticated techniques are used, such as the methods described below.

The Nelson-Siegel/Svensson Extrapolation Methods place lower reliance on the 30-year spot rate for extrapolation purposes, and as a result, the hedging portfolios derived using these methods tend to be highly spread across the range of tradable and observable interest rates.

For the Nelson-Siegel method, the spot curve can be seen as a linear combination of three component functions with different shapes—a flat curve, a sloped curve and a humped curve. The Svensson method is an extension of the Nelson-Siegel model that adds an additional humped curve and allows a more diverse set of yield curves to be modeled. It can potentially introduce a large amount of multicollinearity when fitting against the actual yield curve data.

For any extrapolation and finite set of parameters, the long rate is guaranteed to exist and to be finite; however, it will not remain constant across periods. Use of the observable yield curve is high as the model can be fitted taking into account all relevant data. The Nelson-Siegel standard model has three time-dependent parameters and factors; however, a more generalized version specifies λ (shape parameter) as a time-dependent parameter. Similar comments apply to the Svensson model. The parameters of the model are directly related to tradable quantities. There is little risk of multicollinearity in the Nelson-Siegel model because the various factors all correspond to different types of yield curve movement.

An advantage of a method such as Nelson-Siegel is that it allows one to match the slope of the fitted curve at the start of the extrapolation. This would be an advantage if the slope of the fitted curve was rising relatively rapidly or if it was turning downwards and away from the UFR at its final point.

Although it rests only on estimation of four parameters, the model is able to capture the most common shapes that the yield curve takes on in practice. It is very popular because of the relatively low parameter number and easy computation. In spite of its small parameter number, it can capture most of the yield curve shapes. The disadvantage of the Nelson-Siegel model lies in precision. It is not a good choice if very accurate estimations are required or when one is trying to model a complex yield curve. The four parameters are β_1, β_2, β_3 and τ.

The forward rate curve estimation is

$$f_t(\tau) = \beta_{1,t} + \beta_{2,t} e^{-\tau/\lambda} + \beta_{3,t}(\tau/\lambda) e^{-\tau/\lambda},$$

where

- τ represents the term of the forward rate
- $\beta_{1,t}$, β_{2t}, $\beta_{3,t}$ represent time-dependent stochastic variables and
- λ is a shape parameter.

The Svensson approach proposed an extension of the Nelson-Siegel model by adding an additional hump-shaped element, as shown in Figure 1.

Figure 1

Nelson-Siegel and Nelson-Siegel Svennson Curves

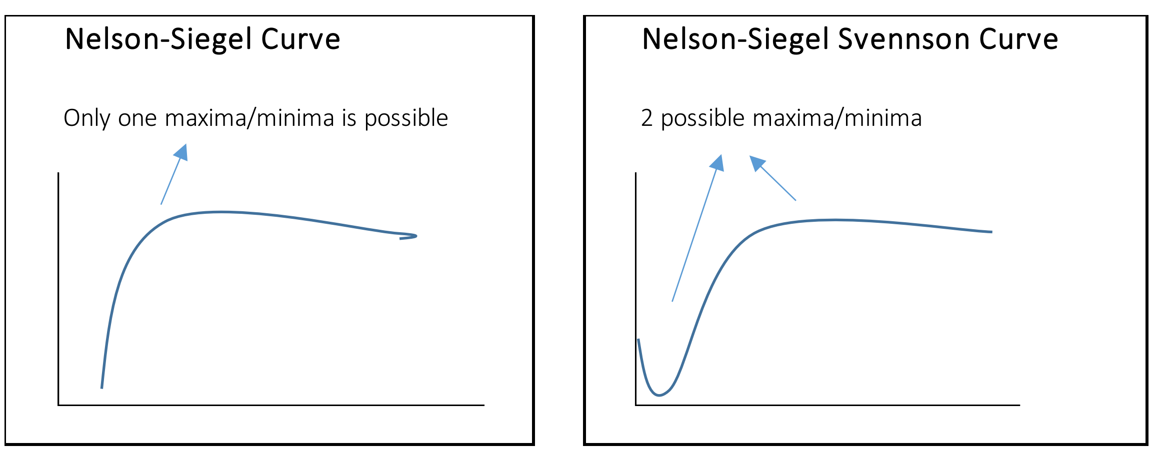

It was intended that this model should be capable of producing a better fit to yield curve shapes with more than one local minimum or maximum. The model for forward rates at a given point in time (t) is

$$f_t(\tau) = \beta_{1,t} + \beta_{2,t}e^{-\tau/\lambda 1} + \beta_{3,t}(\tau/\lambda_1)e^{-\tau/\lambda 1} + \beta_{4,t}(\tau/\lambda_2)e^{-\tau/\lambda 2}$$

where

- τ represents the term of the forward rate
- $\beta_{1,t}$, β_{2t}, $\beta_{3,t}$ $\beta 4_{,t}$ represent parameters that determine the shape of the yield curve and
- λ_1, λ_2 are shape parameters.

The Cairns Extrapolation Method is an exponential-type model similar to that of Nelson-Siegel. It has lower reliance on the 30-year spot rate for extrapolation purposes. As a result, the hedging portfolios derived using this method tend to be more highly spread across the range of tradable and observable interest rates. For any extrapolation and finite set of parameters, the long rate is guaranteed to exist and to be finite; however, it will not remain constant across periods. The usage of the observable yield curve is high as the model can be fitted taking into account all relevant data. The approach does not display intertemporal consistency. The standard model has four time-dependent parameters and factors. As with the Nelson-Siegel/Svensson methods, the parameters of the model are directly related to tradable quantities. There is little risk of multicollinearity because the various factors all correspond to different types of yield curve movement.

For $\tau > M$ the spot curve at time (t) is given by

$$s_t(\tau) = \beta_{0,t}h_{0,t}(\tau) + \beta_{1,t}h_{1,t}(\tau) + \beta_{3,t}h_{3,t}(\tau) + \beta_{4,t}h_{4,t}(\tau)$$

where

$$h_{0,t}(\tau) = 1,$$

$$h_{i,t}(\tau) = (1-e^{c_i\tau})/c_i\tau, \text{ for } i = 1 \text{ to } 4.$$

The Smith-Wilson Extrapolation Method is a class of models where the long forward rate is a fixed input parameter and does not vary over time as bond prices change. It allows the long-term forward rates to converge toward the chosen “infinite” rate and is an almost perfect solution for hedging the long-term interest rate risk. However, for hedging purposes, one weakness with the approach is its absolute reliance on the longest observable spot rate. For any extrapolation, the long rate is guaranteed to exist and remain constant across periods. The usage of the observable yield curve is high as the model can be fitted taking into account all relevant data. The approach does display intertemporal consistency because movements in the parameters of the model can potentially be defined by many intertemporally consistent models. There are potentially many parameters and, therefore, potentially many factors affecting the curve. However, this depends on the process governing changes in the time-dependent parameters. The parameters of the model are directly related to tradable quantities.

The single most common method seen is the Smith-Wilson approach. The main benefit of Smith-Wilson is its simplicity; however, it has a number of drawbacks:

- A single parameter (alpha) controls speed of extrapolation to UFR (ultimate forward rate) and smoothness of curve after market data region. This typically results in a very bumpy forward curve.
- The method works best when there are few market data points—for example, when fitting to swap curves. The method tends to perform poorly when fitting to noisy data, such as when fitting government bond yields (unless there is significant filtering of the data, which introduces subjectivity).

The Smith-Wilson method uses the available data to exactly fit bond prices where data are available and to extrapolate them by using a weighted average of the last observable data point and the predetermined UFR. The speed of convergence to this level is determined by the mean reversion parameter. The Nelson-Siegel model takes a more empirical rather than theoretical approach to yield curve modeling, being able to fit yield curves of diverse shapes (humped, S-shaped, inverted and monotonic).

The Smith-Wilson performs poorly when fitting to a large number of data points; for example, when fitting to bond yields, forward curves are typically very bumpy.

Another important concept is that of the speed of convergence to the UFR, represented by the α parameter in the Smith-Wilson method; this also determines the smoothness of the curve. The higher the α, the bigger the weight of the UFR (hence faster convergence), whereas a lower α gives more weight to the market data.

When calculating the Smith-Wilson yield curve the following specifications apply:

The input parameters are

- The Ultimate Forward Rate (UFR)
- α, the speed of convergence to the UFR

Smith-Wilson assumes that the discount factor, $P(\tau)$, at time t is determined by

$$P_t(\tau) = e^{-f_\infty \tau} + \sum \xi_{i,t} K_i(\tau) \text{ (sum from } i = 1 \text{ to I)},$$

$$K_i(\tau) = \sum c_{ij} W(\tau, u_j) \text{ (sum from } j = 1 \text{ to } J_i),$$

$$W(\tau,u) = e^{-f_\infty(\tau+u)}[\alpha \min\{\tau,u\} - e^{\alpha \max\{\tau,u\}} \sinh(\alpha \min\{\tau,u\})],$$

where

- $c_{i,j}$ represents the jth cash flow on the ith bond used to calibrate the price function, and u_j represents the term of the respective cash flow
- ξ represents a series of time-varying parameters used to fit the actual yield curve
- K represents a set of kernel functions for each input observable bond price and
- W is a symmetric function known as "Wilson's function."

The Cubic Spline Extrapolation Method extends the cubic spline used to fit the market data to the unconditional horizon. It is a class of models in which the long forward rate is a fixed input parameter and does not vary over time as bond prices change. It allows the long-term forward rates to converge toward the chosen "infinite" rate.

A cubic spline is a piecewise cubic polynomial joined at so-called knot points. At each knot point, the polynomials that meet are restricted so that the level and first two derivatives of each cubic are identical. Each additional knot point in the spline adds one independent parameter, because three of the four parameters of the additional cubic polynomial are constrained by the restriction. By increasing the number of knots, cubic splines provide increasingly flexible functional form.

Yield curves cannot be directly observed; instead, they can only be inferred from market prices. As a result, when estimating a yield curve in practice, one attempts to estimate a smooth curve from a relatively small number of fixed-income prices. The cubic spline approximation to a yield curve smooths out the interior bumps by matching both slope and curvature at each interior yield to maturity node that is inferred from some market price. An immediate problem that arises is how to handle the beginning and endpoints in the data set. The natural cubic spline sets the slope at the endpoints to zero.

Although there is no prescribed methodology, a preferred (in-house) method might be a Cubic-Spline Nelson-Siegel (CSNS)—this uses a cubic spline to fit the market data part of the curve and the Nelson-Siegel functional form to handle the extrapolation phase. This has more parametric freedom to achieve good fits to market data and control the extrapolation behavior—although this comes with a slight increase in complexity. The fit to data and extrapolation phases should be considered together to ensure sensible behavior over the entire curve (especially at the transition). A further observation (based on experiences with Solvency II) is that the assumptions that feed into an extrapolation method can have far greater impact than the technical methodology.

Other uses of the CSNS include the following:

- A yield curve fit model that is updated for different dates as needed, using observed government bond yields.
- Slow linear reversion to an assumed long-term rate, and cubic splines for a short-to-medium term. This is robust in terms of estimation and is consistent with views that reversion probably exists, and to the extent that it does, it is (very) slow.

For more formulas relating to the above methods, see the "Long term forecasting and hedging of the South African yield curve" thesis in the References.

3.4 Comparing the Nelson-Siegel, Svensson and Smith-Wilson Methods

In Table 2 we compare some of the more often utilized "complex" methods.

The Nelson-Siegel and Smith-Wilson methods are quite different in the way they are formulated. For extrapolation, the Smith-Wilson method relies on the last known observation (at the Last Liquid Point [LLP]) and on the defined UFR, and the curve is created based on a weighted average of both for the period of convergence. On the other hand, the Nelson-Siegel method uses all the observed data to fit a curve and then uses the factor loadings, or the component coefficients, to extrapolate the remainder of the curve beyond the LLP. The Svennson method is an extension of the Nelson-Siegel where a second medium-term "hump" factor with a separate decay parameter is added.

Table 2

Comparison of Frequently Used "Complex" Methods

	Pros
Nelson-Siegel	The three components give the model enough flexibility to capture monotonic, humped and S-type curves often typically observed in yield curve data
	Parameters are easy to estimate and have simple, intuitive explanations
	Widely used by central banks and practitioners
Svennson	Can more easily fit term structures with more than one local maximum or minimum, thereby allowing for a broader and more complicated range of yield curves
Smith-Wilson	Can be applied to raw market data
	Provides a perfect fit to liquid market data
	Cons
Nelson-Siegel	Highly nonlinear, which has been reported to cause estimation problems
	Cannot handle *all* yield curve shapes
	Assumes forward rates are always positive and the discount factor approaches zero as maturity increases
Svennson	No significant improvement of the estimates when compared with the Nelson-Siegel model
	The model is highly nonlinear, which can make the estimate of the model difficult
	The overparameterization of the model can cause convergence problems
Smith-Wilson	Expert judgment is needed for the choice of alpha (the speed of convergence to the UFR)
	$P(t)$, the discount factor, may become negative

3.5 Smith-Wilson Model

Recently the IAIS proposed an international capital standard, International Financial Reporting Standard 17, in which the calculation involves use of an extended yield curve. The IAIS uses a variant of the Smith-Wilson technique, which is more robust, and many companies are using it. It can be used to illustrate how various IAIS proposals would have worked over the last several decades had they been in place at that time.

Below we give a representation of the results of two Smith-Wilson models with the components, or inputs used, described below Figure 2. The inputs are as follows:

UFR	= Ultimate Forward Rate (long-term equillibrium rate)
α	= measure for the speed of convergence towards the UFR with the Smith-Wilson method (start with the value 0,1)
T2	= number of years until the UFR is reached with the linear extrapolation method
Illiquidity premium quote	= percentage of the illiquidity premium that is to be used (0% or 100% with a binary approach)
Illiquidity premium	= maximum level of the illiquidity premium
Last maturity of illiquidity premium	= number of years until the illiquidity premium should no longer be used
Reduction period of illiquidity premium	= number of years before the last maturity of the illiquidity premium where a linear reduction of the illiquidity premium starts
Maximum spread beween methods at T2	= when the Smith-Wilson and the linear extrapolation methods are calibrated, this is the maximum allowed spread at T2
Calculation date	= date when the calculations are executed

Figure 2 has α equal to 0.05 and the number of years until the UFR is reached of 90.

Figure 2

Comparison of Forward Rate Smith-Wilson and Forward Rate Linear Models, α = 0.05

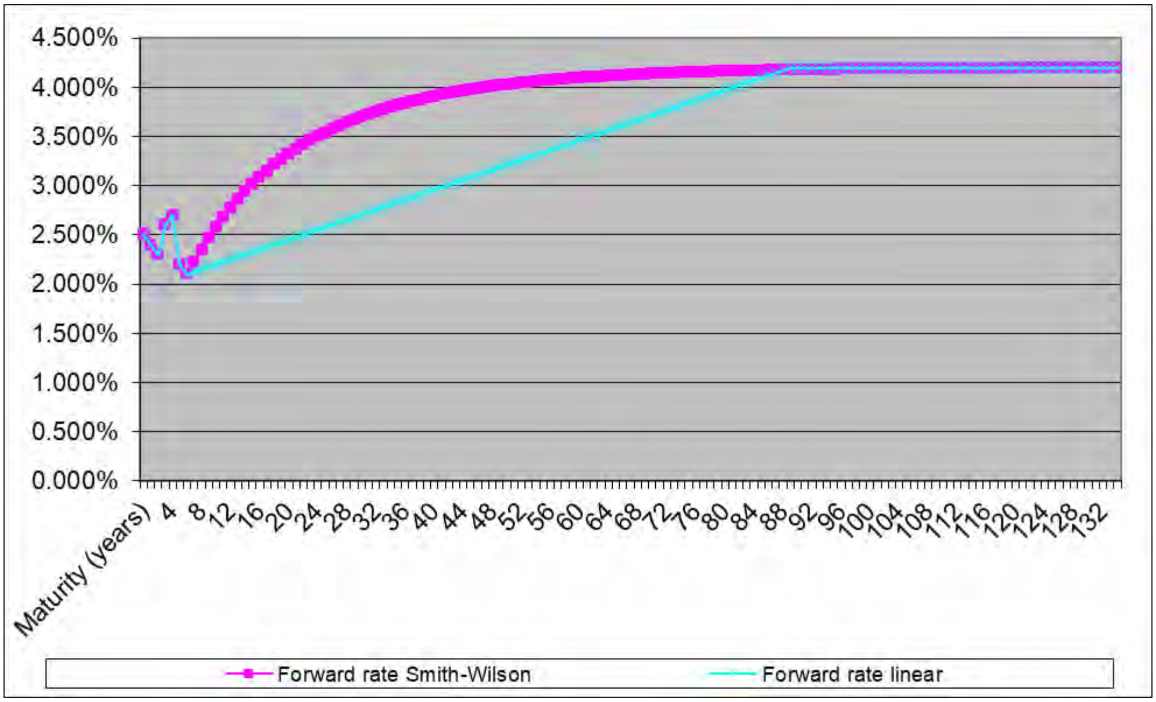

UFR	4.20%
α	0.05
T2 (years)	90
Illiquidity premium quote	100%
Illiquidity premium	0.00%
Last maturity of illiquidity premium (years)	5
Reduction period of illiquidity premium starts at (years)	5
Maximum spread beween methods at T2	0.03%
Calculation date	12/31/2009

Figure 3 has α equal to 0.1 (higher than Figure 2), and the number of years until the UFR is reached is 60 (lower than Figure 2).

Figure 3

Comparison of Forward Rate Smith-Wilson and Forward Rate Linear Models, α = 0.1

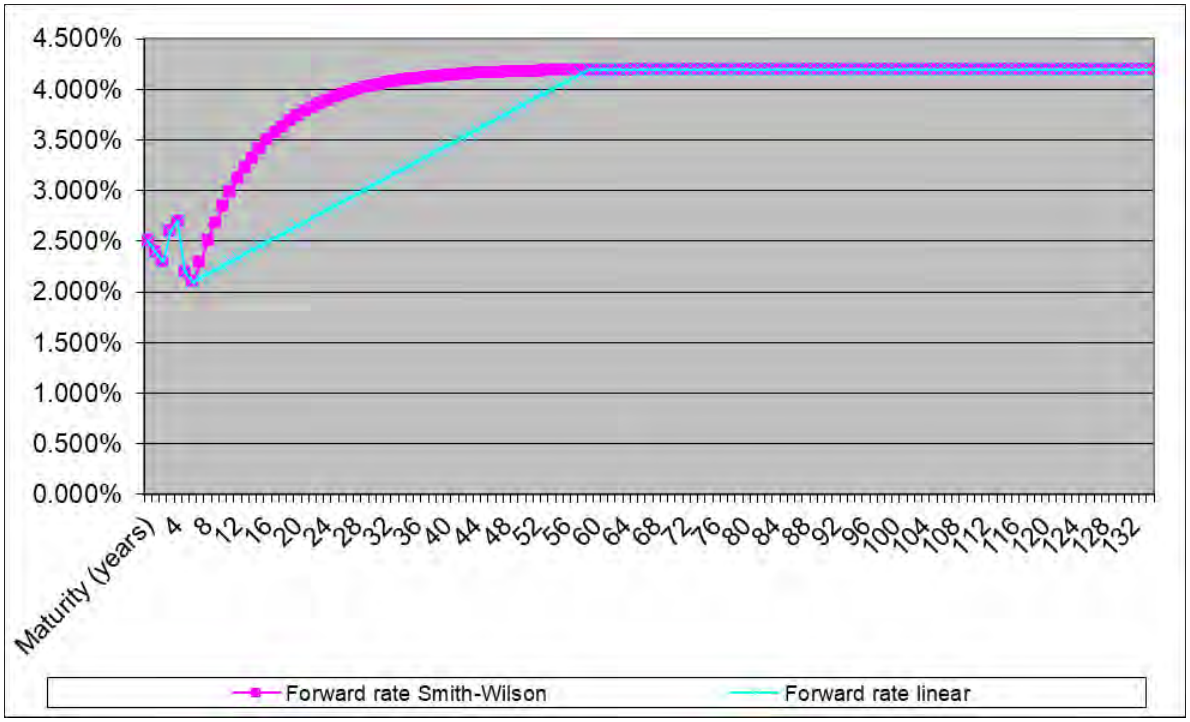

UFR	4.20%
α	0.1
T2 (years)	60
Illiquidity premium quote	100%
Illiquidity premium	0.00%
Last maturity of illiquidity premium (years)	5
Reduction period of illiquidity premium starts at (years)	5
Maximum spread beween methods at T2	0.03%
Calculation date	12/31/2009

A higher α in Figure 3 shows that convergence to the UFR is much quicker over a shorter period.

For more information on the Smith-Wilson calculation, please see the CEIOPS research paper "Solvency II calibration paper" (page 13) in the References.

3.6 Panelist Views on Practical Applications

Various methods are commonly used for yield curve extrapolation.

One could develop a proprietary approach, which forecasts two points on the curve as a multiple of the then-current long-term rate. The long-term rate is set by an autoregressive (Definitions) series. Other points on the curve are by interpolation. The benefits include that it is workable, and results have been validated in a statistical sense against history and current professional calibration requirements. The drawbacks include that there is not much theoretical justification other than the fact that results appear to be statistically valid. Being a proprietary approach, it may not automatically be accepted by clients or stakeholders.

Pricing areas can use primarily simple methods for rate indications on guaranteed annuity products and some others. Any extrapolation method involves significant risk, and since the potential riskiness and/or accuracy and/or bias of various methods is hard to evaluate, it may be appropriate to choose to put resources into evaluating risk of any proposed rate structure, and not put resources into making “better” forecasts.

These days many tend to use Smith-Wilson where mandated and CSNS in other situations. The key benefit of the CSNS method is that, when appropriately parameterized, it allows curves to be produced in a highly automated way while reliably meeting quality criteria:

- Good quality of fit to market data
- Smooth transition between market data and extrapolation phase
- Appropriate convergence to UFR.

A potential limitation is that the parameterization requires some care in the setup, but this is achievable with appropriate attention and expertise.

Another method is a linear reversion to a long-term rate, with the reversion generally starting near the end of the observable market bonds. This approach does not vary by products or purpose. The main benefits are stability and robustness when re-estimating. Other parameterizations (such as Nelsen-Siegel, Smith-Wilson or Merrill Lynch exponential splines) often varied significantly in their speed of reversion, which is undesirable. They are also sometimes harder in enforcing constraints compared to the linear reversion.

The main disadvantages are not restricted to this approach but include the following:

- Plausibility of shape: Most people would expect a decay shape where the speed of reversion slows across the durations.
- Choice of reversion rate: There is some judgment in setting the long-term (reversion) level, compared to a flat assumption.
- Choice of time to achieve full reversion. One panelist chose 50 years, which feels slow enough (faster would not be consistent with their research) but is somewhat ad hoc.

The UFR could be incorporated into the risk-free curve, which is linearly graded in 10 years from the last liquid point at 30 years (to 40 years). The starting point of the extrapolation is the average of the annual forward rates in years 21–30. The UFR is set at management discretion, trying to align with expectations for market rates. It avoids volatility through the averaging in years 21–30, and it avoids discontinuities. Disadvantages are that it relies on some degree of management discretion in setting UFR. For Solvency II, benefits are that it is compliant with the requirements. It is also simple. Drawbacks are that under Solvency II companies are forced to use it. There is only one parameter to control the interpolation and extrapolation, which is less flexible.

Some practical challenges that have arisen from the various applied approaches include the following:

- Many approaches seem to be very complex while still requiring a large amount of judgment and discretion.
- Any method not based on setting the future forward curve can lead to unusual and unrealistic patterns of forward rates.
- Generally, parametric methods (Vasicek, Nelson-Siegel, Svensson) lack the flexibility to accurately fit market data (and extrapolation behavior). Particularly for liability valuation, this is a critical requirement.

- Flat extrapolations are potentially oversimplified (particularly in markets where liabilities are longer that the longest dated traded instrument) and fail to account for many practitioners' preference for a UFR.
- A 50-year discount rate curve (or 200 quarters) makes the curve a little unwieldy for valuations where the long-term rate is less relevant. Some actuaries prefer a simpler approach for their valuations.
- Thinking through the last liquid point can be a challenge. Getting agreement from stakeholders is always a challenge. U.S. versus European expectations can add complexities, because UFR tends to be lower in countries with lower interest rates.

Some drawbacks that led panelists to not use other methods and instead use the one they do are the following:

- Stability in reversion speed was a major factor. They preferred the "shapes" produced by their natural cubic splines; they seemed a bit more faithful to the data compared to other exponential-function-based splines.
- One panelist tried to use the Nelson-Siegel for internal purposes, but it was too complex.

3.7 Panelist Views on Mechanics

The panelists were asked how the following items are determined when extrapolating the yield curve.

For the *Ultimate Long-Term Forward Rate*, including components such as the expected inflation and expected real short-term rate, the rate is usually a combination of judgment and officially published requirements. Generally a simple extrapolation of the current long spot and forward rates is used. Using each of the components, these are estimated from pooled (across multiple currencies) historical data. Consulting with other long-term rates (such as the government's intergenerational reports) as well as historical data on long-dated bonds is helpful to check for reasonableness. The UFR accounts for expectations of long-term real interest rate and inflation. Term premia and convexity adjustment are not included in the determination of the UFR.

For the *Duration of the Ultimate Long-Term Forward Rate,* one panelist stated lacking anything definitive, using 20 years is a reasonable approach. Another stated they use 30 years and then grade over another 30 years. Some comments were more general such as "as needed for product pricing application and as long as needed for projection." Others go as long as 50 years, which seems plausible based on looking at other countries with longer-dated bonds (U.K./Canada/U.S.). The duration is driven by the last liquid point for market data and a reasonable convergence period.

The Speed of Convergence to the Ultimate Forward Rate is based on judgment and historical data, and it could be defined by a simple method. If based on volatility of the resulting interest rates, it typically results in a slower reversion than Solvency II curves. One panelist stated the convergence is linear from around 15 years through to year 50.

The shape and smoothness of the transition from the observed rates to the extrapolated rate generated by the algorithm might be defined by a simple method. An important consideration is the smoothness of the transition. Some prefer a smoother transition as compared to Solvency II. Others argue this is wrong and should jump from discontinuity to smooth. Linear interpolation is a popular transition despite being slightly nonintuitive compared to a decay curve. Integrating this component with natural cubic splines for the short-medium term estimates has the advantage of a seamless single-estimation procedure.

When fitting these prices to estimate the forward rates, fitting the prices exactly results in a very irregular (bumpy) looking forward curve. The reason for this is that that bonds of adjacent maturities can trade at prices that are slightly off-trend as traded volumes across different maturities vary. In other words, liquidity effects can lead to noise in reported bond prices. Because of this noise, a better estimate of the true underlying forward curve is obtained by imposing some smoothness on the estimated forward rates. In fact, various studies have shown that smooth forward rate curves perform better than bumpy ones in predicting the prices of bonds deliberately left out of the forward rate estimation process; that is, the smoother curves give better out-of-sample performance.

The mechanics or processes used to fit the curve include least squares with some judgment, automated processes, implemented as a solver optimization in Excel, linear programming and an interpolation method called monotome convex, which ensures that the continuous forward rates are positive.

Other relevant considerations include the credit quality associated with the yield curve. If one starts by building a curve for sovereign or "risk-free" investments and adds a curve of credit spreads, then a method is needed to extend the curve of credit spreads. The IAIS, in its field study, proposed a method where the credit spread in the ultimate forward rate is zero, although that was done for conservatism and may not be viewed as realistic. Also, as seen and experienced in recent times, negative rates can happen. Any acceptable model and target system should be able to handle negative rates and reinvestment risk.

3.8 Panelist Views on Complexity

Panelists were asked how simple or complex do they consider their extrapolation methodologies.

Many state that their method is simple and adequate. They believe that other methods involve more complex math without much evidence that they are actually any more theoretically justifiable. A simple approach is used because the liabilities involve nonguaranteed elements, which minimizes the effect of any choice of approach as long as minimum guarantees are not in the money. Others prefer a complex method, perhaps using a CSNS method that is at the more sophisticated end of the spectrum. Although they may have investigated many other techniques, they feel this is the best solution given their requirements for highly accurate fits and an automated production process. Although the heart of the estimation is not too complex (a single fit of around a six parameter function applied to observed bond yields), to do this requires a fair bit of setup (e.g., splitting up the bonds into a stream of zero-coupon bonds or applying some duration-based weights). When deciding, they were unconvinced that any other more complex shapes offered practical advantages. One panelist stated they endeavor to use the simplest model possible, but no more simple than is necessary to be consistent with the market and economic principles.

3.9 Panelist Views on Hedgeability

Panelists were asked if they consider hedgeability of the interest rate risk in picking their extrapolation methodologies. By definition, interest rates beyond the investible universe are not in general entirely hedgeable, so the question reduces to residual risk estimation of various hedge strategies. For example, this risk could be estimated based on traditional duration or partial-duration based hedging strategies, reflecting the extent to which the hedge portfolio offsets changes in the estimated price of liabilities. Alternatively, given the proposed extrapolation of pricing rates, current investable universe and a stochastic rate generator, this risk could be defined in terms of the surplus distributional results over 30–50 years, and the extent to which these results can be improved by varying investment strategy and/or hedging approaches.

Some panelists stated they hedge to a degree. This manifests as using as much of the available market data as possible, often resulting in a longer LLP than the Solvency II approach. Some panelists stated that hedgeability is considered, although they generally provide reserve estimates rather than implement hedging strategies. A flat assumption appears poorer for hedging purposes. Another panelist does not currently consider hedgeability; however, they are working toward a dynamic programming approach to hedge design. This will entail modeling future markets and hedges in a consistent way.

3.10 Panelist Views on Software Selection

The main consideration in the selection of the software package(s) to be used to extrapolate the yield curve is using a vendor or in-house model. Most of the panelists indicated they use in-house software. One built their own "'Yield Curve Builder" software, which includes the CSNS, Smith-Wilson and a bootstrap technique. This is integrated into their broader service delivery infrastructure to allow automated production and integration with market data feeds. Another panelist stated that Excel was sufficient for their purpose and has the advantage of being easy to incorporate into the workflows for most people in the office. The simplicity of the approach allows for an in-house software package that is similar to an Excel-based approach (but the software package allows for more robustness in auditing). As for vendor packages, one panelist currently uses MATLAB with Financial and Financial Instruments toolboxes. This is useful but requires doing it yourself. They are evaluating third-party ESGs presently.

3.11 Panelist Views on Governance and Responsibilities

In many cases, panelists indicated that the actuary ultimately prepares the recommendation of extrapolated yields. This could be the pricing actuary or the actuary that prepares all cash-flow-testing assumptions. Some other responses were as follows:

- The product/pricing group makes the recommendation; the asset liability management (ALM)/investment strategy area does the evaluation. The result is approved by senior management based on the risk assessment.
- Curves are produced using automated infrastructure and managed by the modeling operation team. Input assumptions are set by the research team and updated on a two-year cycle. Assumptions are peer reviewed by the research team. Produced curves are subject to analyst and senior analyst review before publication.
- The research team prepares recommendations and presents them to the senior actuarial staff. Ultimately, all actuaries are entitled to use whatever assumptions they want for their clients. They encourage (and generally see) a consistent approach across the company.
- The Group Risk Capital Committee (GRCC) sets methodology, while the group risk area performs the calculations. Approval is through the GRCC, but they do not approve the numbers (just the methodology). There are controls to ensure that the numbers are consistent with the governance approval from the committee.
- For risk-neutral analysis, the lead strategist and the head of variable annuity ALM develops the approach. For real-world analysis the corporate economist, who works within the investments division, develops the approach with assistance from strategists in ALM.
- One company has an Assumptions Governance Committee that reviews and approves all economic assumptions with potential material impact on economic, regulatory or accounting results.
- Curves are generated in valuation models. The methodology was developed internally and approved by external audit. Valuation is done daily, so it is generated in the model.

The responses varied as to whether the process is documented under either Model Audit Rule (MAR), Sarbanes–Oxley (SOX) or elsewhere. One panelist noted that the process, methodologies and assumptions are fully documented and governed by their own Data Policy and Calibration Policy. The process is controlled and documented sufficiently to satisfy internal standards and, when applicable, external requirements such as MAR or SOX.

3.12 Additional Panelist Comments

Panelists were asked how projected negative interest rates affect their model. One panelist stated that their models did not produce negative rates, since at that time they believed they were impossible. This was before the 2008 crash. Some permit negative rates and think this is an important consideration. Their models could handle observed negative yields and reflect them in the fit. Others have no impact from negative rates. They use a normal, rather than lognormal, volatility model, to help avoid problems with negative rates.

In an ESG with a sound underlying theoretical model (no arbitrage, consistent interrelationship between benchmark rates and prices), negative rates occurring on paths in some trials are acceptable. When it comes to expected or mean reversion target rates, negative levels can reasonably occur and be acceptable in early periods, but negative rates at and beyond extreme dates are not acceptable. They are not generally supported by reasonable economic expectations and tend to arise as artifacts of the approach (for instance, using splines to extrapolate instantaneous forwards).

Regarding how their models handle negative interest rates, many panelists stated that the model has a floor. Some models are built to avoid them. Others have models that can handle observed negative yields and reflect them; however, users can choose to apply a floor if required. In the realistic model based on shifted lognormal there are ultimate floors but no floors in market-consistent models.

When asked what models are missing from our list, panelists responded the following:

- Vasicek functional form—a type of one-factor short rate model that describes interest rate movements as driven by only one source of market risk
- Merrill Lynch exponential spline method—a parametric model that specifies a functional form for the discount function
- Minimal Fisher information—a variational principle that, when applied with the proper constraints needed to reproduce empirically known expectation values, determines the best probability distribution that characterizes the system
- Graeme West Monotome Convex—this has a mechanism that ensures that the forward rates generated by the method are positive (equivalently, that the discount factor curve is decreasing).

In any application, the fundamental issue is risk assessment. According to one respondent, many actuaries invest too much time in sophisticated extrapolation methods, which provide false confidence in the accuracy of the result.

The applicability of the methods to areas beyond liabilities in the insurance industry, such as assets, pension and other post-retirement benefits, could be used in any application where the fundamental issue is risk assessment. This could be applicable anywhere when evaluating long-dated cash flows. Each method should be just as appropriate for application to any long horizon analysis of assets and liabilities as it is for liabilities in the insurance industry.

The methods are applicable in jurisdictions beyond the U.S., because the math should work anywhere. Some jurisdictions are less stable, so results might be so volatile as to not be useful. They may be useful in large, mature, stable economies with liquid rate market instruments. These sorts of approaches would not be useful in economies subject to substantial event risk, illiquid or government-controlled markets, frequent or likely economic regime change, inflation volatility and the like.

One panelist stated that the two things they learned about yield curve extension (particularly when doing economic capital work immediately postcrisis) are (1) there isn't any "right" answer and (2) given point 1, ease of explanation and simplicity (and consistency with markets with long observable rates) can be more important than theoretical "purity."

Section 4: Concluding Remarks

It is important to recognize that these extrapolation methods are models, and both the models and the assumptions going into the models need to be strongly vetted by the user to ensure applicability of the model and the appropriateness of the assumptions for the purpose for which it is used. One must determine through the assumptions if they will model an average or extreme view. A company wanting to be more conservative may model with more extreme down assumptions, whereas another that wants stability may use more average assumptions.

There is a wide range of modeling methods, from simple linear models to more complex spline models. A good extrapolation model strikes the right balance—practicality on the one hand, together with the ability to capture the most important attributes and most critical features of history.

Based on the research, it appears that many of the methods used in practice are similar to those discussed and analyzed in the theoretical literature. The assumptions that feed into an extrapolation method may have a greater impact than the technical methodology. However, the choice of the method does have an impact on the results. Using the current year forward rate extrapolated out into the future will have a much different result than a method that grades over time.

A large number of experts seem to be using the simpler models, favoring simplicity over complex models. Models that are too simple, however, can miss the true risks and may not appropriately capture tail events. At the other extreme, a good model does not “overfit” the data, reducing the ability to produce simulations beyond the historical data itself.

A good extrapolation model will produce results that are relevant to historical facts. A common tendency is to overweight the importance of the recent past. The danger in placing too much focus on recent risks is that one can forget that, over long periods, the economy can move rates to new and different places. With a longer-term horizon, one must avoid the temptation to influence a view.

We took a closer look at a few of the more popular methods, each having their own pros and cons. In selecting an extrapolation method, one needs to determine the best fit for their particular needs. If using one for Solvency II, the Smith-Wilson method is required. In other circumstances where curve flexibility and easy to understand parameters are needed, perhaps the Nelson-Siegel method would be preferred.

Section 5: Expert Panelists

Special thanks goes to those who took the considerable time and effort required to provide thoughtful and detail-oriented responses to the questionnaire.

Name	Company
Stephen J. Strommen , FSA, MAAA, CERA	Blufftop, LLC
Joseph Koltisko, FSA, MAAA, CFA	New York Life
Andrew M. Erman, FSA, MAAA	Transamerica/Aegon
Robert Reitano, FSA, MAAA, CERA, PhD	Brandeis University
Peter Mulquiney, FIAA, PhD	Taylor Fry
Hugh Miller, FIAA, PhD	Taylor Fry
Ricky Power, FSA, FIA, CERA	Moody's Corporation
Steven Prince, FSA, FCIA, MAAA	RSM Canada
Joseph D'Anna, PhD, CFA	Brighthouse Financial
John Manistre, FSA, FCIA, MAAA, CERA, PhD	Retired
Josh Windsor, FSA, FIA, MAAA	NAIC
Vadim Zinkovsky, FSA	Metlife

Section 6: Acknowledgments

The researchers' deepest gratitude goes to those without whose efforts this project could not have come to fruition: the Project Oversight Group and others for their diligent work overseeing questionnaire development, analyzing and discussing respondent answers, and reviewing and editing this report for accuracy and relevance.

Project Oversight Group Members:

Robert Reitano, FSA, MAAA, CERA

Alex Ngo, FSA, FCIA, CERA

Andrew Harris, ASA

Linda Peach, ASA, MAAA

Nathan Deboer, FSA, MAAA

Peter Mason, FIA

Larry Rubin, FSA, MAAA, FCA, CERA

James Kosinski, FSA, MAAA

At the Society of Actuaries:

Max Rudolph, FSA, MAAA, CERA—Society of Actuaries/External Project Manager

Steven Siegel, ASA, MAAA

Barbara Scott, Sr. Research Administrator

References

The following documents have been reviewed to prepare portions of the information provided in this report (along with links, if available).

- Antonio, David, and David Roseburgh. 2010. Fitting the yield curve: Cubic spline interpolation and smooth extrapolation. Version 3.0. https://www.scribd.com/document/283751499/Fitting-the-Yield-Curve-and-Cubic-Spline-Interpolation.
- Barrie, A., and J. Hibbert. 2008. A framework for estimating and extrapolating the term structure of interest rates. Exposure draft. http://janroman.dhis.org/finance/Curves%20-%20Interpolation/A_Framework_for_Estimating_and_Extrapolating_the_Term_Structure.pdf.
- CEIOPS (Committee of European Insurance and Occupational Pensions Supervisors). 2010. Solvency II calibration paper. CEIOPS-SEC-40-10. https://eiopa.europa.eu/ceiops-archive/documents/advices/ceiops-calibration-paper-solvency-ii.pdf.
- EIOPA (European Insurance and Occupational Pensions Authority). Tool for extrapolation risk free curves. 2009. https://eiopa.europa.eu/Publications/QIS/ceiops-tool-extrapolator-risk-free-rates_en.xls.
- Manistre, B. John. 2009. A risk management tool for long liabilities: The static control model. Society of Actuaries. https://www.soa.org/essays-monographs/2009-erm-symposium/mono-2009-m-as09-1-manistre.pdf.
- Mulquiney, Peter, and Hugh Miller. 2012. A topic of interest—How to extrapolate the yield curve. Institute of Actuaries of Australia, General Insurance Seminar: Tides of Change, November 12–13, 2012, https://www.actuaries.asn.au/Library/Events/GIS/2012/GIS2012PaperMulquineyMiller.pdf, or "Comment" in *Actuaries,* August 2013, https://www.actuaries.asn.au/Library/AAArticles/2013/ActuariesAUG2013p35-37.pdf.
- Pedersen, Hal, Mary Pat Campbell, Stephan L. Christiansen, Samuel H. Cox, Daniel Finn, Ken Griffin, Nigel Hooker, Matthew Lightwood, Stephen M. Sonlin, and Chris Suchar. Economic scenario generators: A practical guide. 2016. Society of Actuaries. https://www.soa.org/Files/Research/Projects/research-2016-economic-scenario-generators.pdf.
- Thomas, Michael Patrick. Long term forecasting and hedging of the South African yield curve. Thesis, University of Pretoria. https://repository.up.ac.za/bitstream/handle/2263/25595/dissertation.pdf.
- Vuong, Cindy. 2013. Yield curve models for extrapolation. Towers Watson. https://prezi.com/f-efenw8ktkj/yield-curve-models-for-extrapolation/.

Appendix A: Questionnaire

Project Scope

Investigate methodologies that are employed for valuing liability cash flows that extend beyond the maximum observable portion of the yield curve. This will be accomplished using research approaches that include reviewing technical specifics of appropriate literature along with interviews of individual experts regarding current practices. The results of these approaches will be analyzed and summarized in a final report.

General Notes on Questionnaire

Thank you for participating in this questionnaire. Please note the following:

1. If a particular question does not apply to you, please state "Not Applicable".
2. Your responses are confidential. The final research report will be based on anonymized responses and would not involve specific information on any particular company or client.
3. To the extent possible, for each answer you provide, please include the applicable country(ies), industry(ies), product(s), intended use of the extrapolation, etc.
4. By "Yield Curve Extrapolation Methods", we mean methods in place to extend the yield curve beyond the current investable universe for purposes of determining the yield rates that can be used to determine the current value of cash flows occurring later than the current investable universe.
5. We prepared a brief summary of yield curve extrapolation methods. Please see Appendix A of this questionnaire for details.
6. This survey is intended to be updated electronically within this MS Word document. The "Response" section for each question will expand as responses are typed in.

Interviewee Background and General Experience

1. Please share examples of your experience in extrapolating the yield curve for purposes of valuing long term liabilities (for example, for purposes of internal capital calculations)
    a. Would you consider your experience to be theoretical, practical, or both (please explain)?
    b. Does your experience include the development and use of economic scenario generators (please explain)?
    c. In what range of liability categories have you applied the extrapolation of the yield curve? Potential examples include life insurance, annuities, health insurance, Property & Casualty and pension liabilities.
    d. Does your experience include developing and/or implementing regulatory requirements that involve yield curve extrapolation, such as Solvency II (please explain)?

General Information on Yield Curve Extrapolation Methods

2. Please discuss what you consider to be industry approaches for extrapolating the yield curve, and describe specific situations (for example, specific products and specific applications) in which each is used. Please also discuss what you would consider benefits and drawbacks of the various approaches you described.

3. Please share regulatory approaches or actuarial guidelines that you work with for the extrapolation of the yield curve, including the rationale for the choices (for example, IFRS, Solvency II, etc.).

Practical Information on Yield Curve Extrapolation Methods

4. What product(s) do you use the yield curve extrapolation for?
5. What purpose(s) do you use yield curve extrapolation for—liability products, asset projections only, valuation, pricing, regulatory, etc.? Please indicate if your responses vary by products.
6. What method(s) do you use for yield curve extrapolation and why? Please indicate if your responses vary by products or purpose.
    a. What are benefits and drawbacks of your chosen method?
    b. Please share examples of practical challenges that have arisen from the various applied approaches.
    c. What are drawbacks that led you to not use other methods and instead use the one you do?
7. How are the following items determined when you extrapolate the yield curve?
    a. The Ultimate Long-Term Forward Rate, including components such as the expected inflation, expected real short-term rate, term premia, and convexity adjustment
    b. The Duration of the Ultimate Long-Term Forward Rate
    c. The Speed of Convergence to the Ultimate Forward Rate, if applicable to your method
    d. Shape and smoothness of the transition from the observed rates to the extrapolated rate
    e. What mechanics or processes do you use to fit the curve?
    f. Any other relevant item depending on your chosen method
8. How simple or complex do you consider your extrapolation methodologies?
    a. If your approach is a complex approach, please help us understand why that method was considered better than alternative or simpler methods. What value is gained from using the more complex method? Have you performed an analysis to determine that the complex method provides better results than a simple method?
    b. If your approach is a simple approach, please help us understand why that was used. Is it due to immateriality of the product to the business, or ease of use?
    c. If your approach is a simple approach, have you performed an analysis to quantify the benefits associated with using a complex method? If so, did minimum benefits play a role in your use of the simple approach?
9. Do you consider hedgeability of the interest rate risk in picking your extrapolation methodologies? What considerations are in place?
10. What criteria and features do you consider in the selection of the software package(s) to be used to extrapolate the yield curve?
    a. What do you consider to be the benefits and drawbacks of the software chosen, compared to other options available?
    b. What was lacking in the softwares that were not chosen?
11. Please share any other relevant information regarding yield curve extrapolation used at your Company or for your clients that we may not have specifically requested.

Governance and Standard of Practice

12. Who ultimately prepares the recommendation of extrapolated yields?
13. Who ultimately approves the extrapolated yields?
14. Is this process documented under either MAR, SOX, or elsewhere?
15. Negative interest rates:
    a. How do projected negative interest rates affect your model?
    b. How does your model handle negative interest rates? Are they overridden, or is there a floor on the rates used?

Other Questions on Yield Curve Extrapolation Methods

16. The following methods are methods that are generally used for yield curve extrapolation:
    - Simple extrapolation method
    - Flat Rate extrapolation method
    - Linear First Order extrapolation method
    - Other First Order extrapolation method
    - Nelson-Siegel extrapolation method
    - Svensson extrapolation method
    - Cairns extrapolation method
    - Smith/Wilson extrapolation method
    - Cubic spline extrapolation method

    Please respond to the following questions on these extrapolation methods based on your experience or research.

    a. What methods are missing from the list above?
    b. Which of these methods have you used or investigated?
    c. What do you consider to be the benefits or drawbacks to any of these extrapolation methods?
    d. Please provide specific examples of practical challenges that you may have observed from the use of specific methods.
    e. Please share the applicability of the methods to areas beyond liabilities in the insurance industry, such as assets, pension, and other post-retirement benefits.
    f. Please comment on the applicability of these methods in jurisdictions beyond the US.

Appendix B: Definitions

An **autoregressive model** is when a value from a time series is regressed on previous values from that same time series.

Convexity is a measure of the curvature in the relationship between bond prices and bond yields that demonstrates how the duration of a bond changes as the interest rate changes (second derivative).

Forward rate is a term used in both bond and currency trading to represent the current expectations of future bond interest rates or currency exchange rates. In bond trading, the forward rate is calculated based on interest rates for various maturities. These rates are typically plotted on a graph as a yield curve.

Intertemporal consistency means that the activities that an individual plans now to carry out in the future are the activities that he or she actually carries out when the future arrives.

The method of **least squares** is a standard approach in regression analysis to approximate the solution of overdetermined systems, that is, sets of equations in which there are more equations than unknowns. The **linear least-squares** problem occurs in statistical regression analysis; it has a closed-form solution.

Long-term rates are typically higher than short-term rates. The yield curve is the chart that appears when you graph interest rates compared to maturity or length of the loan.

In statistics, **multicollinearity** (also called **collinearity**) is a phenomenon in which one predictor variable in a multiple regression model can be linearly predicted from the others with a substantial degree of accuracy.

A **power relation(ship)** corresponds to a formula with a power in it. You can use these formulas for lines that increase very fast.

First-order has an order of one, especially denoting mathematical equations involving only the first power of the independent variable or only the first derivative of a function.

Spot rate is the price quoted for immediate settlement on a commodity, a security or a currency. The spot rate, also called the "spot price," is based on the value of an asset at the moment of the quote. As a result, spot rates change frequently and sometimes dramatically.

About The Society of Actuaries

The Society of Actuaries (SOA), formed in 1949, is one of the largest actuarial professional organizations in the world dedicated to serving 32,000 actuarial members and the public in the United States, Canada and worldwide. In line with the SOA Vision Statement, actuaries act as business leaders who develop and use mathematical models to measure and manage risk in support of financial security for individuals, organizations and the public.

The SOA supports actuaries and advances knowledge through research and education. As part of its work, the SOA seeks to inform public policy development and public understanding through research. The SOA aspires to be a trusted source of objective, data-driven research and analysis with an actuarial perspective for its members, industry, policymakers and the public. This distinct perspective comes from the SOA as an association of actuaries, who have a rigorous formal education and direct experience as practitioners as they perform applied research. The SOA also welcomes the opportunity to partner with other organizations in our work where appropriate.

The SOA has a history of working with public policymakers and regulators in developing historical experience studies and projection techniques as well as individual reports on health care, retirement and other topics. The SOA's research is intended to aid the work of policymakers and regulators and follow certain core principles:

Objectivity: The SOA's research informs and provides analysis that can be relied upon by other individuals or organizations involved in public policy discussions. The SOA does not take advocacy positions or lobby specific policy proposals.

Quality: The SOA aspires to the highest ethical and quality standards in all of its research and analysis. Our research process is overseen by experienced actuaries and non-actuaries from a range of industry sectors and organizations. A rigorous peer-review process ensures the quality and integrity of our work.

Relevance: The SOA provides timely research on public policy issues. Our research advances actuarial knowledge while providing critical insights on key policy issues, and thereby provides value to stakeholders and decision makers.

Quantification: The SOA leverages the diverse skill sets of actuaries to provide research and findings that are driven by the best available data and methods. Actuaries use detailed modeling to analyze financial risk and provide distinct insight and quantification. Further, actuarial standards require transparency and the disclosure of the assumptions and analytic approach underlying the work.

Society of Actuaries
475 N. Martingale Road, Suite 600
Schaumburg, Illinois 60173
www.SOA.org