

Methodologies for Valuing Cash Flows That Extend Beyond the Yield Curve

By Benjamin Leiser and Jack Kerbeshian

In 2018, Risk & Regulatory Consulting conducted a study on methodologies used for yield curve extrapolation to value liability cash flows that extend beyond the maximum observable portion of the yield curve. The study was sponsored by the Society of Actuaries (SOA) Committee on Finance Research. We performed research on the methods available in theory and used in practice and also developed and provided questionnaires to a broad group of subject-matter experts with strong industry representation to comment on these methods, in order to provide a comprehensive view of the yield curve extrapolation methods. We developed the interview questionnaire based on an initial review of the literature and covered topics such as

- 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; and
- practical challenges that arise from various methods.

We then supplemented the initial research with results of the survey and interviews, including any theoretical and practical issues noted with the methods. Both the research and the survey included details of how these methods are applied, as well as observations on the benefits, drawbacks and prevalence of their use. This article provides a summary of our study. The complete report, including a summary of panelists' views, can be found on the SOA's website (www.soa.org/resources/research-reports/2019/yield-curve-report/).

OVERVIEW

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 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.

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.

EXTRAPOLATION

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 and reliable market instrument, and, second, extrapolating from the longest available and reliable 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:

1. What is an appropriate assumption for the infinite-maturity, unconditional forward rate of interest?
2. 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.

A good yield curve estimation method must deliver extrapolated curves that are credible at some point in time and where changes over time in extrapolated rates can be justified.

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, which are the additional returns an investor may expect as compensation for the longer-term investment and are represented by 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:
 - **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 (not to be confused with credit or equity risk premia).
 - **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 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.

SAMPLE 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. In this section we list several methodologies along with some detail on each method.

The Simple Extrapolation Method

The simple extrapolation method is simple to implement. It has two variations:

- **The simple monopole method.** This method assumes a constant single forward rate for all durations greater than 30.
- **The simple dipole method.** This variation uses the maximum observable (often 30-year) forward rate beyond that point.

The Flat Rate Extrapolation Method

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 non-observable 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, and while this is based on a tradable quantity, it could be limited when liquidity is low.

The Linear First-Order Extrapolation Method

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 gradient (slope of rates) and scale (level of rates). If the two factors are determined exclusively from the observable yield curve, then they will be hedgeable. This method assumes that the forward rates beyond M years follow a first-order linear progression of the form $f_{+}(\tau) = a + b \times \tau, \tau > M$,

where

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

Other First-Order Extrapolation Methods

Two more first-order extrapolation methods bear discussing:

- **The power spot rate extrapolation method.** This model assumes that forward rates beyond the longest observable spot rate follow a power relation. This method assumes that the forward rates beyond M years follow a power progression of the form

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

- **The exponential spot rate extrapolation method.** This method assumes that forward rates beyond the longest observable spot rate follow an exponential relation. It assumes that the forward rates beyond M years follow an exponential progression of the form

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

The Nelson-Siegel-Svensson Extrapolation Methods 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

For the Nelson-Siegel method, the spot curve is expressed 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.

The forward rate curve estimation is

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

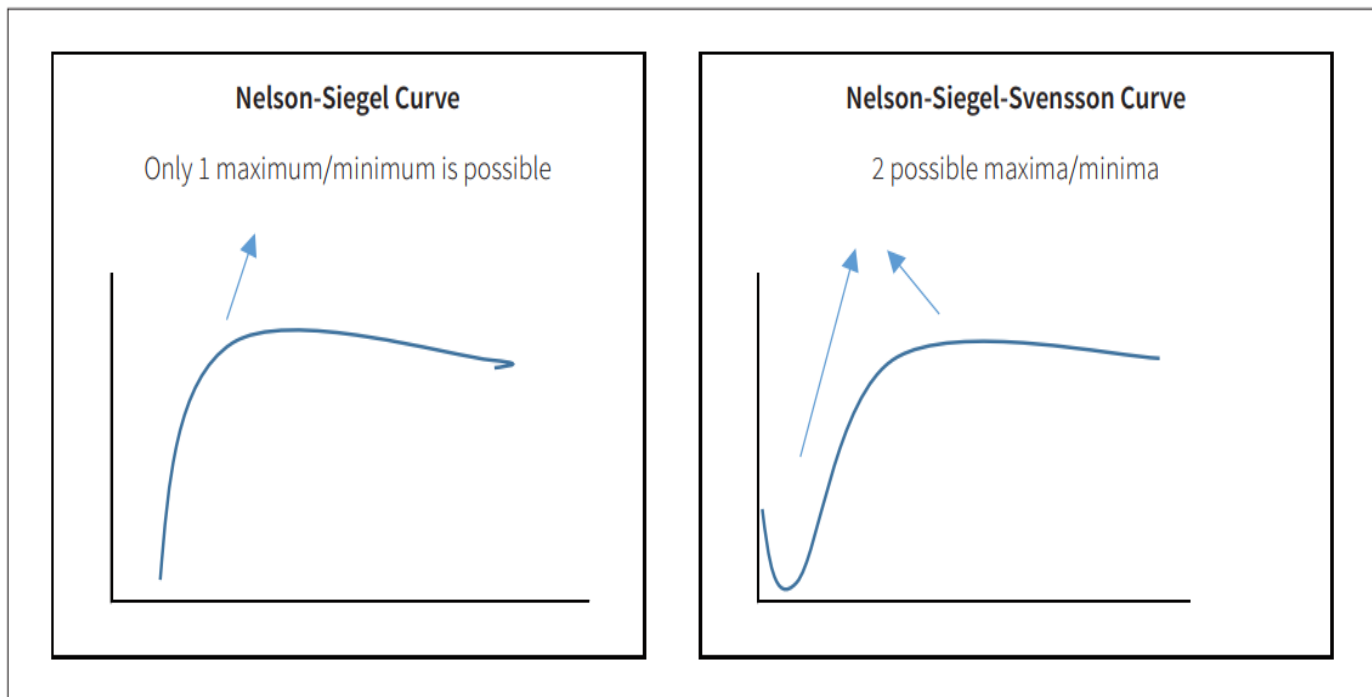
to be highly spread across the range of tradable and observable interest rates.

where

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

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

Figure 1
Comparing Nelson-Siegel Curve with Nelson-Siegel-Svensson Curve



The Smith-Wilson extrapolation method

The Smith-Wilson extrapolation method 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 and provides a strong basis for hedging the long-term interest rate risk.

The input parameters are

- the UFR, and
- α , the speed of convergence to the UFR.

$$P_i(t) = e^{-f^t} + \sum \xi_{i,t} K_i(t) \text{ (sum from } i = 1 \text{ to } I),$$

$$K_i(t) = \sum c_j W(t, u_j) \text{ (sum from } j = 1 \text{ to } J),$$

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

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

- c_{ij} represents the j th cash flow on the i th 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

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.

COMPARING THE NELSON-SIEGEL, SVENSSON AND SMITH-WILSON METHODS

In Table 1, 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. For its part, 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 Svensson method is an extension of the Nelson-Siegel in which a second medium-term “hump” factor with a separate decay parameter is added.

Table 1
Comparison of Frequently Used “Complex” Methods

Model	Pros	Cons
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	Highly nonlinear, which has been reported to cause estimation problems
	Parameters are easy to estimate and have simple, intuitive explanations	Cannot handle all yield curve shapes
	Widely used by central banks and practitioners	Assumes forward rates are always positive and the discount factor approaches zero as maturity increases
Svensson	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	No significant improvement of the estimates when compared with the Nelson-Siegel model Highly nonlinear, which can make the estimate of the model difficult Overparameterization of the model can cause convergence problems
Smith-Wilson	Can be applied to raw market data	Requires expert judgment for the choice of alpha (the speed of convergence to the ultimate forward rate)
	Provides a perfect fit to liquid market data	$P(t)$, the discount factor, may become negative

SAMPLE OF PANELIST VIEWS

Expert industry panelists were given questionnaires to comment on methods used to extrapolate the yield curve both in practice and in theory. We include here some of the responses they provided that helped to supplement our research with respect to the various methods that are commonly used for yield curve extrapolation.

These days many tend to use Smith-Wilson where mandated and cubic spline Nelson-Siegel (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; and
- 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.

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 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 than 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.

When appropriately parameterized, the CSNS method allows curves to be produced in a highly automated way while reliably meeting quality criteria.

Any extrapolation method involves significant risk, and because the potential riskiness, accuracy and bias of various methods is hard to evaluate, it may be appropriate to choose to put more resources into evaluating the risk of any proposed rate structure than into attempting to make “better” forecasts.

- For the UFR, 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 UFR, 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 when looking at 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 UFR is based on judgment and historical data, and it could be defined by a simple method. One panelist stated the convergence is linear from around 15 years through to year 50.

- The speed of convergence to the UFR is based on judgment and historical data, and it could be defined by a simple method. 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, while others indicate the transition should jump from discontinuity to smooth. Linear interpolation is a popular transition despite being slightly nonintuitive compared to a decay curve.
- 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 monotone convex, which ensures that the continuous forward rates are positive.

Many of the panelists seem to agree that there isn't any "right" answer and, therefore, ease of explanation, simplicity and consistency with markets with long observable rates can be more important than theoretical "purity."

CONCLUSION

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 if the assumptions and model result in 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, with the ability to capture the most important attributes and most critical features of history on the other.

Based on the research and the survey of industry experts, 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 as great, or at times even a greater impact than the technical methodology. However, the choice of the method itself does have an impact on the results. Using the current year forward rate

extrapolated out into the future will have a much different result from 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.

Through our research and the survey as described in this article, we took a closer look at a few of the more popular methods, each having their own pros and cons. We were not surprised to find that in selecting an extrapolation method, users must determine the best fit for their particular needs. 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 any more theoretically justifiable. As 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.

We would like to thank the SOA and the project oversight group for their contributions and support throughout this process.



Benjamin Leiser, FSA, MAAA, is a consulting actuary at Risk & Regulatory Consulting. He can be reached at ben.leiser@riskreg.com.



Jack Kerbeshian, FSA, MAAA, is an actuarial consultant at Risk & Regulatory Consulting. He can be reached at jack.kerbeshian@riskreg.com.