Yieldcurves: Yield Curve Fitting, Analysis, and Decomposition in R

by Charles Coverdale

Abstract The yieldcurves package provides an end-to-end toolkit for term-structure analysis in R. It fits parametric yield curves via the Nelson-Siegel (1987) and Svensson (1994) models and interpolates non-parametrically via cubic splines. From a fitted curve it extracts spot rates, instantaneous forward rates, discount factors, and par rates via bootstrap stripping. It computes Macaulay and modified duration, convexity, Z-spread, and key rate durations for zero-coupon and coupon bonds. It decomposes yield-curve movements into principal components following Litterman and Scheinkman (1991) and into level, slope, and curvature factors. It reports carry and roll-down over a user-specified horizon and a menu of slope and butterfly measures. Every function accepts plain numeric inputs and has no dependencies beyond base R and cli. The package is available on CRAN with source on GitHub.

1 Introduction

The **yieldcurves** package implements the standard toolkit of term-structure analysis: parametric and non-parametric curve fitting, extraction of forward rates and discount factors, risk measurement, and factor-style decomposition of curve movements. Nineteen exported functions share a uniform `yc_` prefix and a curve-object-based interface, so a workflow that begins with `yc_nelson_siegel()` can chain through `yc_forward()`, `yc_duration()`, `yc_carry()`, and `yc_pca()` without reshaping or re-coding.

The state of R infrastructure for yield-curve work has been thin. **YieldCurve** (Guirreri, 2022) fits Nelson-Siegel and Svensson models but requires xts or zoo objects, does not compute duration, convexity, Z-spread, or any factor decomposition, and has not been updated since 2022. The more comprehensive **termstrc** was archived from CRAN in 2018 after repeated failures against the current CRAN check machinery and has not been resurrected. The **fBonds** package exposes bond-pricing primitives but not curve estimation. Neither existing package covers the full workflow from raw bond quotes through to carry, roll-down, key rate durations, and principal-component decomposition.

yieldcurves consolidates these operations into a single package with no heavy dependencies. Imports are **cli** (error messages) plus base R `graphics` and `stats`. R 4.1.0 or later is required. The only suggested package is **testthat** for the test suite, which contains over one hundred and seventy tests across every exported function. All inputs are plain numeric vectors: maturities in years, rates as decimals. Outputs are S3 objects with `print()` and, where geometric intuition is central, `plot()` methods.

2 Background

Yield-curve analysis combines three families of techniques: smoothing across the maturity dimension, risk-measure derivation from a smooth curve, and factor-style decomposition of curve time series.

Parametric smoothing. Nelson and Siegel (1987) proposed a four-parameter functional form consisting of a level component, a slope component decaying exponentially with maturity, and a curvature component producing a medium-maturity hump. Over two decades later the Nelson-Siegel form remains the workhorse of central-bank yield-curve estimation, used by the Federal Reserve, the European Central Bank, the Bank of Canada, and over twenty other central banks in the Bank for International Settlements survey (Bank for International Settlements, 2005).

Svensson (1994) generalises Nelson-Siegel with a second curvature term, giving a six-parameter form with enough flexibility to match the occasionally-bi-humped US Treasury curve. Diebold and Li (2006) showed that the Nelson-Siegel parameters can be interpreted as dynamic latent factors (level, slope, curvature), reframing the curve as a vector autoregression in parameter space and providing a natural forecasting model; Diebold and Rudebusch (2013) provide the textbook treatment. Gürkaynak et al. (2007) publish a real-time Svensson curve for the US Treasury market that has become the standard research dataset for term-structure studies. Bauer and Rudebusch (2020) document that the long-run component of that curve, corresponding to Nelson-Siegel's β_0, has drifted noticeably downward over the past four decades as the perceived level of the long-run neutral rate has fallen.

Non-parametric smoothing. Cubic splines (McCulloch, 1971) and tension splines (Fisher et al., 1995) avoid the tight functional restriction of Nelson-Siegel. They can fit any observed curve exactly at the cost of potentially wiggly forward rates. **yieldcurves** exposes `stats::splinefun()` cubic splines for users who want non-parametric interpolation alongside the parametric fits.

Risk measures. Macaulay (1938) defined duration as the value-weighted average time to a bond's cash flows; modified duration, the sensitivity of bond price to a parallel shift in yields, is the Macaulay quantity divided by $(1 + y/k)$ (Hicks, 1939). Convexity is the second derivative. The Z-spread is the constant additive shift to the zero-coupon curve that prices a bond at par. Key rate durations (Ho, 1992) decompose a bond's overall duration into exposures to movements at specific points on the curve. These measures are the inputs to portfolio hedging and scenario analysis and must be computed against a smoothed curve rather than a raw set of quoted yields.

Factor decomposition. Litterman and Scheinkman (1991) showed that the first three principal components of US Treasury yield changes explain over 97 per cent of daily variation and have robust interpretations as level, slope, and curvature factors. A thirty-year panel of monthly curves still reproduces this decomposition almost exactly. The 2s10s and 2s30s spreads and the 2s-5s-10s butterfly are the daily curve-shape indicators traders and central-bank economists watch. More formally, Diebold and Li (2006) reframe Nelson-Siegel as a dynamic factor model; Piazzesi (2010) surveys the arbitrage-free affine class that embeds these factors in a no-arbitrage pricing kernel.

Carry and roll-down. Holding a bond over a short horizon yields an expected return equal to the bond's yield (carry) plus the gain from being revalued at the lower yield appropriate to its shorter remaining maturity (roll-down), provided the curve itself does not shift. Carry and roll-down analysis decomposes this expected return by maturity and is a staple of rates-relative-value trading.

3 Package design

yieldcurves is pure R with no compiled code. Runtime imports are **cli**, `graphics`, and `stats`. Nineteen functions cover fitting, extraction, risk measures, and decomposition; the test suite contains over one hundred and seventy tests.

The yc_curve class

Every fitting function (`yc_nelson_siegel`, `yc_svensson`, `yc_cubic_spline`, and the unifying `yc_fit`) returns an object of class `yc_curve` with a common shape: maturities and rates as supplied, the fitted rates at the same maturities, parameter estimates (the `params` list), residuals, method identifier, rate type (`"zero"`, `"par"`, or `"forward"`), and an observation date. Downstream functions dispatch on this class. A user needs to learn one object to use every capability in the package.

Extraction and risk functions

`yc_predict()` returns fitted rates at an arbitrary maturity vector. `yc_forward()` and `yc_discount()` return forward rates and discount factors. `yc_duration()`, `yc_bond_duration()`, `yc_zspread()`, and `yc_key_rate_duration()` compute risk measures. Compounding convention (continuous, annual, or semi-annual) is configurable where relevant.

Decomposition and portfolio analytics

`yc_pca()` runs a principal-component decomposition on a matrix of observed curves (rows are dates, columns are maturities). `yc_level_slope_curvature()` returns the three standard factors defined as fixed linear combinations of points on the curve. `yc_slope()` returns the 2s10s, 2s30s, and 2s-5s-10s butterfly spreads. `yc_carry()` reports carry and roll-down at a user-specified horizon.

Reproducibility

Every numerical result is deterministic given the input. The only stochastic element in the fitting pipeline is the multi-start optimiser for the Svensson model, which accepts a `seed` argument for reproducibility. The package ships no bundled data; the examples in this paper use US Treasury yields obtained from the Federal Reserve Bank of St. Louis Economic Data (FRED) service.

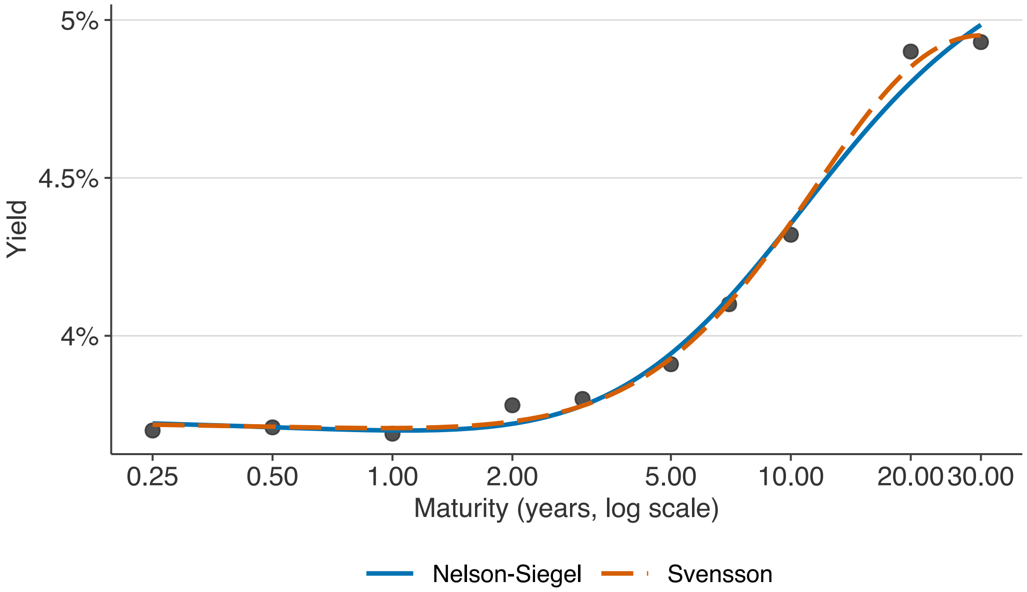

Figure 1: Nelson-Siegel and Svensson fits to the US Treasury yield curve, most recent trading day. Points are the Constant Maturity Treasury rates from FRED at ten maturities. Solid line: Nelson-Siegel, RMSE 4.5 basis points. Dashed: Svensson, RMSE 2.9 basis points. Maturity on a log scale.

Table 1: Parameter estimates for Nelson-Siegel and Svensson fits on the most recent US Treasury curve. Rates are in decimal units. The Nelson-Siegel level β_0 is the asymptotic long-run rate. The negative slope β_1 reflects the upward-sloping curve at the observation date. RMSE in basis points.

Parameter	Nelson-Siegel	Svensson
β_0	0.0536	0.0455
β_1	-0.0162	-0.0082
β_2	-0.0210	0.0578
β_3	—	-0.0459
τ_1	—	8.0000
τ_2	—	5.0000
RMSE (bps)	4.5	2.9

4 Curve fitting

`yc_nelson_siegel()` and `yc_svensson()` estimate the Nelson-Siegel and Svensson parameters by non-linear least squares, taking `maturities`, `rates`, and optional observation `weights`; Nelson-Siegel additionally exposes a `tau` argument. The optimisers use multiple starting values for the decay parameters (τ for Nelson-Siegel, τ_1 and τ_2 for Svensson) and return the best fit by root-mean-squared error. Observation weights are supported for users who wish to down-weight short-end or long-end points.

Figure 1 fits both models to the US Treasury curve on the most recent trading day. Svensson's extra curvature term delivers a roughly one and a half basis point improvement in RMSE at the cost of two additional parameters. The gain is most visible at the very long end, where the two models disagree by about three basis points.

Table 1 reports the parameter estimates at the same date. The Nelson-Siegel β_0 is the asymptotic long-run rate, close to the thirty-year yield. $\beta_0 + \beta_1$ is the short rate. Svensson adds β_3 and τ_2 to capture a second hump when one is present.

`yc_cubic_spline()` fits a cubic interpolating spline via `stats::splinefun()` with the `"natural"` boundary condition. It passes through every input point exactly and is useful for users who have clean data and want no smoothing bias, at the cost of sometimes-wiggly forward rates.

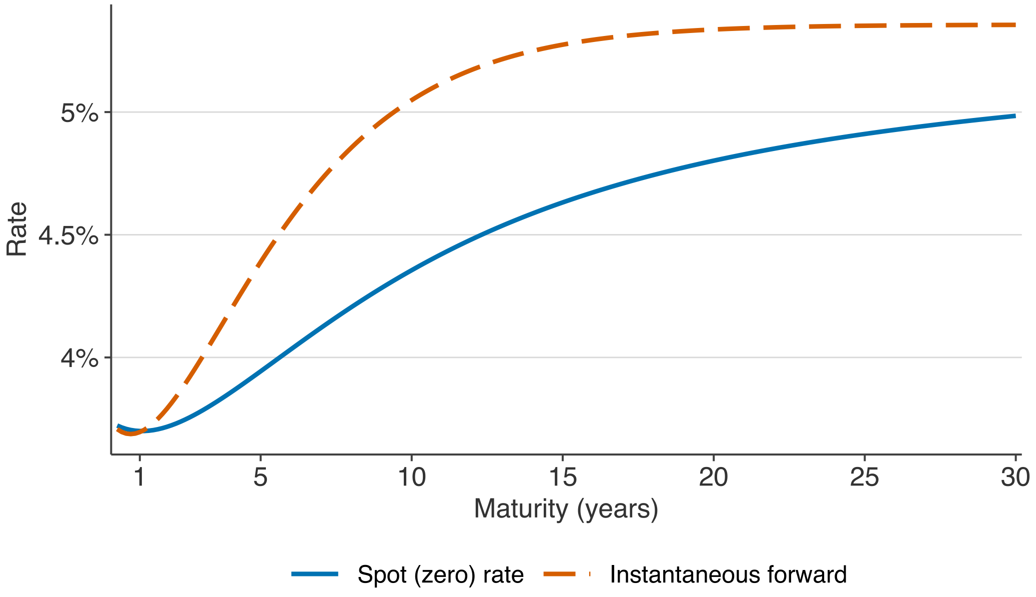

Figure 2: Spot and instantaneous forward rates from a Nelson-Siegel fit to the most recent US Treasury curve. Spot (zero) rate in blue, instantaneous forward in red dashed. The forward curve lies above the spot curve wherever the spot curve is upward-sloping, with the gap widening at maturities where the spot rate rises most steeply.

5 Derived quantities and risk measures

`yc_forward()` returns the instantaneous forward rate at each maturity, derived analytically for Nelson-Siegel and Svensson fits and numerically from the first derivative of log-discount factor for spline fits. `yc_discount()` returns discount factors under continuous, annual, or semi-annual compounding. `yc_par_to_zero()` and `yc_zero_to_par()` convert between par and zero-coupon rates by bootstrap stripping.

Figure 2 plots the spot and instantaneous forward curves from the current Nelson-Siegel fit. Where the spot curve rises, the forward curve sits above it; the gap is the additional yield an investor demands for locking in a longer rate today.

`yc_duration()` returns Macaulay and modified duration for a zero-coupon bond at each maturity, using the curve's own yields. `yc_bond_duration()` does the same for a coupon bond, taking face value, coupon rate, maturity, yield, coupon frequency, and compounding as arguments, and additionally returns convexity. Figure 3 shows the standard textbook result: zero-coupon duration equals maturity, but a coupon bond of the same maturity has lower duration because the intermediate coupons carry weight that shortens the effective payment horizon.

`yc_zspread()` returns the Z-spread for a coupon bond: the constant additive shift to the zero-coupon curve that prices the bond at its market price. It is the standard spread measure for non-Treasury bonds.

`yc_key_rate_duration()` decomposes total duration into exposures to shifts at specific curve points (typically 1, 2, 5, 10, and 30 years) via the `key_rates` argument. These decompositions are the input to key-rate hedging and to scenario analysis that goes beyond parallel shifts.

6 Decomposition

`yc_pca()` performs a principal-component decomposition of a panel of curves. Figure 4 shows the loadings from PCA on the US Treasury curve across monthly observations from January 2019 through the most recent month. The three dominant factors are nearly orthogonal, produce the classical level-slope-curvature pattern documented in Litterman and Scheinkman (1991), and explain 97 per cent, 3 per cent, and 1 per cent of variation respectively, cumulating to 99.9 per cent. The default `scale = FALSE` runs PCA on the covariance of yield changes, so level-heavy samples like this one assign more variance to PC1 than long-run figures in Litterman and Scheinkman (1991); users studying shorter, rangebound samples may prefer `scale = TRUE`.

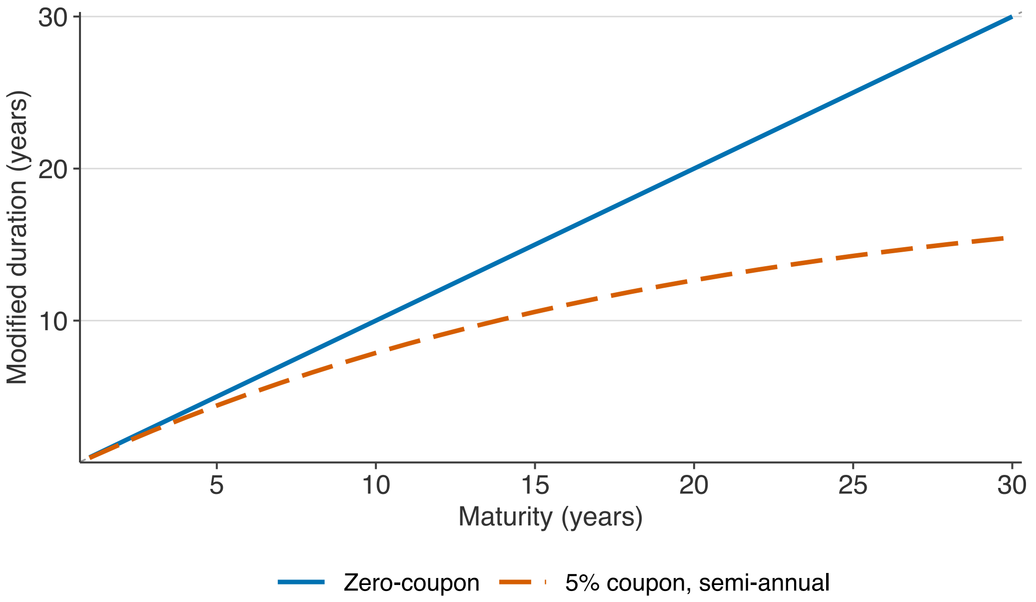

Figure 3: Modified duration versus maturity for zero-coupon and 5 per cent semi-annual coupon bonds. Zero-coupon duration (solid blue) equals maturity minus a small yield-dependent adjustment. A 5 per cent semi-annual coupon bond (red dashed) has substantially lower duration than a zero-coupon bond of the same maturity because the coupon cash flows carry weight at intermediate dates. Yields at each maturity read from a Nelson-Siegel fit to the current US Treasury curve.

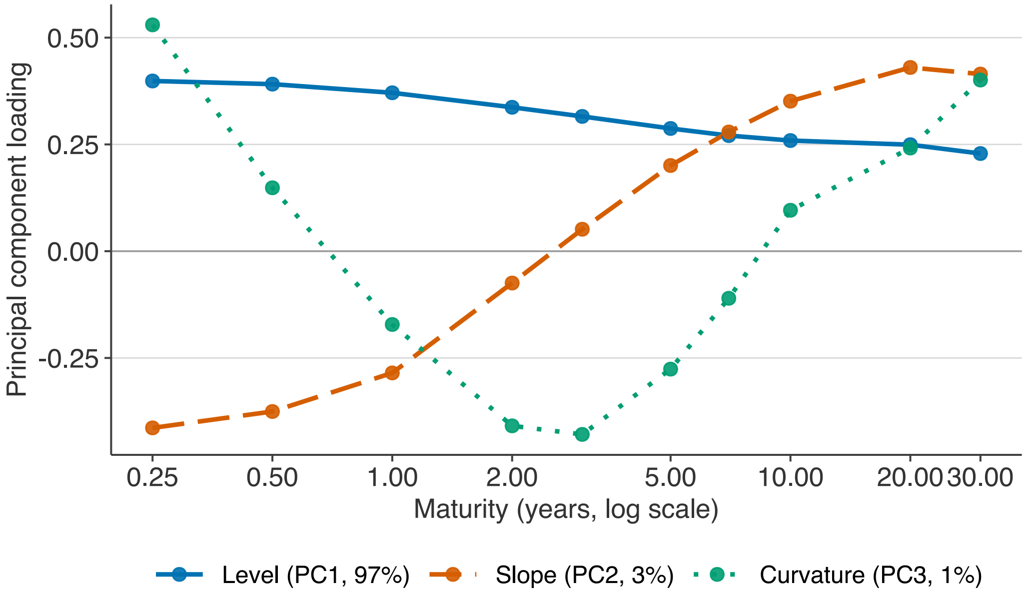

Figure 4: First three principal-component loadings of monthly US Treasury curve changes, 2019 to current. PC1 (level, blue) is roughly flat across maturities, indicating a parallel shift. PC2 (slope, red) is monotonically increasing with maturity, flipping sign around the 3- to 5-year maturity. PC3 (curvature, green) is humped with peaks at the short and long ends and a trough in the intermediate maturities. The three components together explain 99.9 per cent of variation.

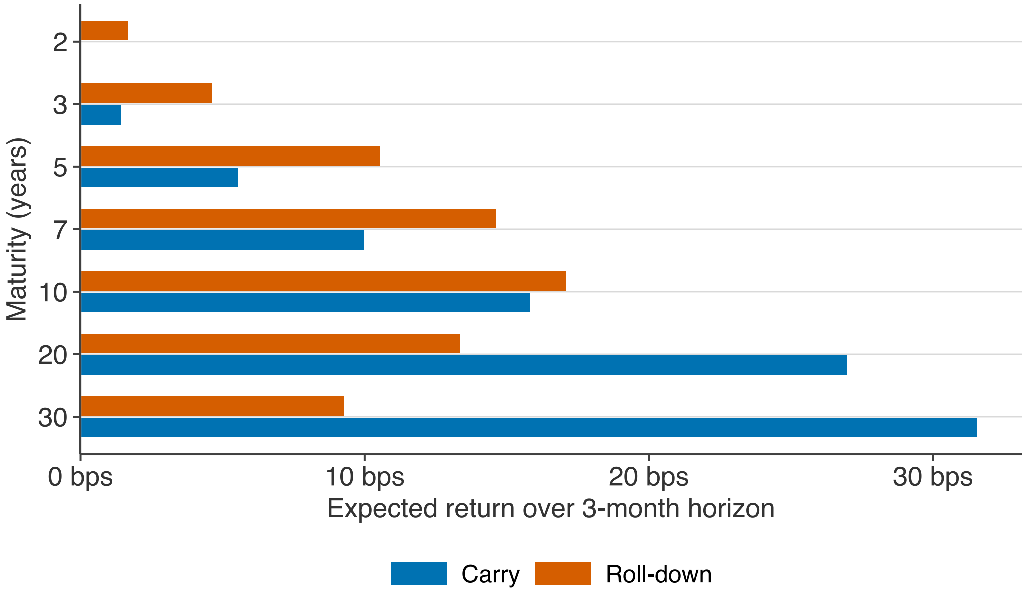

Figure 5: Expected three-month return decomposed into carry and roll-down, US Treasury curve at the most recent date. Bars are in basis points. Carry (blue) is the bond yield less the short-end funding rate (here the 3-month Treasury yield), scaled to the three-month horizon. Roll-down (red) is the capital gain from being revalued at the shorter-maturity yield implied by the current curve, assuming the curve does not shift. The 30-year point has the largest carry because the curve is upward-sloping, but modest roll-down because the curve is nearly flat at the long end.

`yc_level_slope_curvature()` returns the three factors as fixed linear combinations (level is the average, slope is long-short, curvature is a butterfly around the belly), which are easier to compare across dates than rolling PCA.

`yc_slope()` returns the 2s10s (ten-year minus two-year yield), 2s30s (thirty minus two), and the 2s-5s-10s butterfly (twice the five-year minus the two- and ten-year); these three are the standard curve-shape statistics reported in every fixed-income briefing.

7 Carry and roll-down

`yc_carry()` returns the expected return over a specified horizon from holding a bond at each maturity, assuming the curve does not move. The `funding_rate` argument can override the default, which is the curve's own yield at the horizon point. Carry is the bond's yield (less any funding rate if the position is leveraged). Roll-down is the capital gain from being repriced at a shorter maturity's (typically lower) yield. Their sum is the all-in expected return conditional on a static curve.

Figure 5 shows the breakdown across maturities at a three-month horizon, using the current US Treasury Nelson-Siegel fit. The pattern is informative: at the two- to five-year range, carry and roll-down are modest. At the belly (five- to ten-year) they are comparable and additive. At the long end, carry dominates and roll-down falls because the curve flattens.

8 Replication

A canonical yield-curve workflow is five lines:

```
library(yieldcurves)
mats  <- c(0.25, 0.5, 1, 2, 5, 10, 30)
rates <- c(0.0495, 0.048, 0.046, 0.042, 0.041, 0.045, 0.050)
fit   <- yc_nelson_siegel(mats, rates)
yc_forward(fit, maturities = c(2, 5, 10))
yc_carry(fit, maturities = c(2, 5, 10), horizon = 0.25)
```

The same idiom, with `yc_nelson_siegel()` replaced by `yc_svensson()` or `yc_cubic_spline()`,

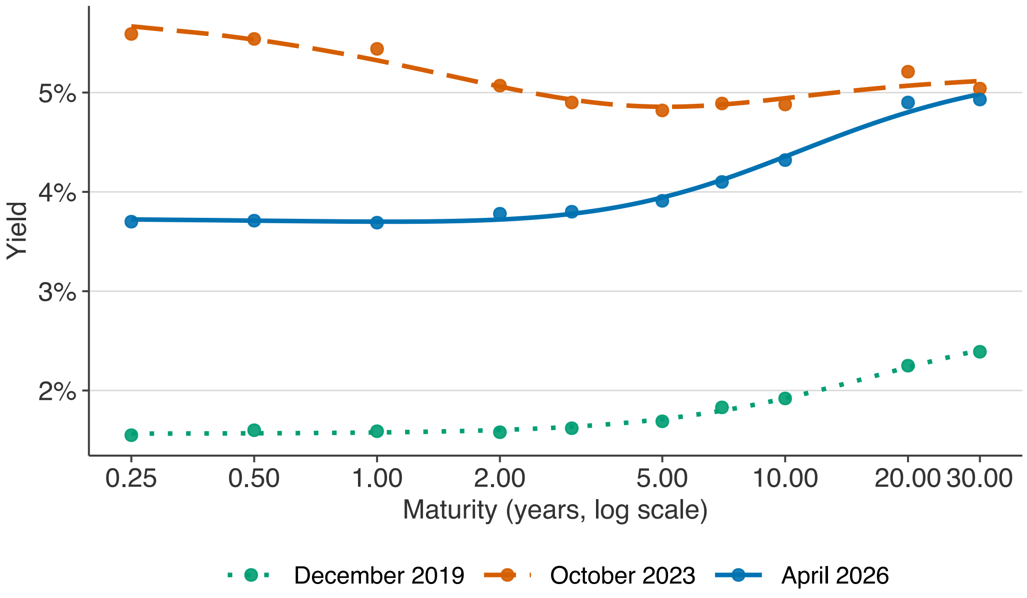

Figure 6: US Treasury yield curves at three dates: pre-pandemic, rate peak, and most recent. Points are FRED Constant Maturity Treasury rates at ten maturities. Lines are Nelson-Siegel fits. December 2019 shows a flat low-rate curve. October 2023 shows an inverted curve around the 2- to 5-year belly. The most recent curve has re-steepened as front-end rates have fallen from their peak.

works for every fitting method. Downstream functions dispatch on the `yc_curve` class and accept plain numeric vectors otherwise.

9 A pandemic-cycle case study

Figure 6 shows the US Treasury curve at three dates spanning the pandemic monetary cycle: the end of 2019, the rate-peak month of October 2023, and the most recent trading day. The pre-pandemic curve is flat and low, reflecting the near-zero policy rate and modest term premia of the prior decade. The October 2023 curve is inverted, with two-year yields above long yields, reflecting the Federal Reserve's tightening cycle and the market's expectation of eventual rate cuts. The current curve is positively-sloped again, with short-end yields having come down from their peak.

Each of these three curves can be passed to the full downstream workflow: `yc_pca()` on the full monthly panel produces the decomposition in Figure 4; `yc_carry()` on the current fit produces Figure 5; `yc_duration()` on any fit produces Figure 3.

10 Limitations

Five limitations apply.

1. **yieldcurves** fits curves at a single point in time. Dynamic term-structure models, including affine Gaussian models following Ang and Piazzesi (2003) and arbitrage-free Nelson-Siegel models, are out of scope. For those, use specialised packages or bespoke estimation.
2. The package does not implement the return-forecasting regressions of Cochrane and Piazzesi (2005) or other predictive exercises using the fitted curve as a regressor. Those are natural extensions but belong in a forecasting-focused package.
3. Par-to-zero bootstrap stripping supports annual and semi-annual frequencies only. Exotic coupon schedules and amortising instruments require user-side preprocessing.
4. The Svensson optimiser is non-convex and can return local minima in pathological cases. The multi-start implementation reduces but does not eliminate this risk; users concerned about convergence should inspect residuals and consider alternative starting values.
5. Real-time curve estimation at a central-bank's publication frequency is not provided. Gürkaynak et al. (2007) publish a continuously-updated Svensson curve for the US Treasury market that users can ingest as an alternative to fitting their own.

11 Conclusion

Yield-curve analysis is a routine part of fixed-income research, portfolio management, and central-bank monitoring, and had been under-served on CRAN. **yieldcurves** consolidates the standard toolkit: Nelson-Siegel, Svensson, and spline fitting; forward rates, discount factors, and par-zero conversions; duration, convexity, Z-spread, and key rate durations; principal-component, level-slope-curvature, and slope-and-butterfly decompositions; carry and roll-down. Every function accepts plain numeric inputs and returns S3 objects with print and plot methods. Planned additions include arbitrage-free Nelson-Siegel following Christensen et al. (2011), a reduced-form term-premium decomposition via `yc_term_premium()` following the regression-based approach of Adrian et al. (2013), return-forecasting helpers for the Cochrane-Piazzesi factor, a Bloomberg-style ASW (asset-swap) spread computation, and optional integration with real-time Svensson curves from central-bank publications. Contributions and bug reports are welcome on GitHub.

Acknowledgements

The Federal Reserve Bank of St. Louis publishes the Economic Data (FRED) service used for the US Treasury Constant Maturity Treasury examples in this paper. The author thanks the R Core Team and CRAN maintainers for the infrastructure that supports open-source statistical software.

Bibliography

T. Adrian, R. K. Crump, and E. Moench. Pricing the term structure with linear regressions. *Journal of Financial Economics*, 110(1):110–138, 2013. doi: 10.1016/j.jfineco.2013.04.009. [p8]

A. Ang and M. Piazzesi. A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables. *Journal of Monetary Economics*, 50(4):745–787, 2003. doi: 10.1016/S0304-3932(03)00032-1. [p7]

Bank for International Settlements. Zero-coupon yield curves: Technical documentation. BIS Papers 25, Bank for International Settlements, 2005. [p1]

M. D. Bauer and G. D. Rudebusch. Interest rates under falling stars. *American Economic Review*, 110(5): 1316–1354, 2020. doi: 10.1257/aer.20171822. [p1]

J. H. E. Christensen, F. X. Diebold, and G. D. Rudebusch. The affine arbitrage-free class of Nelson-Siegel term structure models. *Journal of Econometrics*, 164(1):4–20, 2011. doi: 10.1016/j.jeconom.2011.02.011. [p8]

J. H. Cochrane and M. Piazzesi. Bond risk premia. *American Economic Review*, 95(1):138–160, 2005. doi: 10.1257/0002828053828581. [p7]

F. X. Diebold and C. Li. Forecasting the term structure of government bond yields. *Journal of Econometrics*, 130(2):337–364, 2006. doi: 10.1016/j.jeconom.2005.03.005. [p1, 2]

F. X. Diebold and G. D. Rudebusch. *Yield Curve Modeling and Forecasting: The Dynamic Nelson-Siegel Approach*. Princeton University Press, Princeton, 2013. [p1]

M. Fisher, D. Nychka, and D. Zervos. Fitting the term structure of interest rates with smoothing splines. FEDS Working Paper 95-1, Board of Governors of the Federal Reserve System, 1995. [p1]

S. S. Guirreri. *YieldCurve: Modelling and Estimation of the Yield Curve*, 2022. R package version 5.1. [p1]

R. S. Gürkaynak, B. Sack, and J. H. Wright. The U.S. Treasury yield curve: 1961 to the present. *Journal of Monetary Economics*, 54(8):2291–2304, 2007. doi: 10.1016/j.jmoneco.2007.06.029. [p1, 7]

J. R. Hicks. *Value and Capital: An Inquiry into Some Fundamental Principles of Economic Theory*. Clarendon Press, Oxford, 1939. [p2]

T. S. Y. Ho. Key rate durations: Measures of interest rate risks. *Journal of Fixed Income*, 2(2):29–44, 1992. doi: 10.3905/jfi.1992.408049. [p2]

R. Litterman and J. Scheinkman. Common factors affecting bond returns. *Journal of Fixed Income*, 1(1): 54–61, 1991. doi: 10.3905/jfi.1991.692347. [p2, 4]

F. R. Macaulay. *Some Theoretical Problems Suggested by the Movements of Interest Rates, Bond Yields and Stock Prices in the United States since 1856*. National Bureau of Economic Research, New York, 1938. [p2]

J. H. McCulloch. Measuring the term structure of interest rates. *Journal of Business*, 44(1):19–31, 1971. doi: 10.1086/295329. [p1]

C. R. Nelson and A. F. Siegel. Parsimonious modeling of yield curves. *Journal of Business*, 60(4):473–489, 1987. doi: 10.1086/296409. [p1]

M. Piazzesi. Affine term structure models. In Y. Aït-Sahalia and L. P. Hansen, editors, *Handbook of Financial Econometrics: Tools and Techniques*, volume 1, pages 691–766. Elsevier, 2010. doi: 10.1016/B978-0-444-50897-3.50015-8. [p2]

L. E. O. Svensson. Estimating and interpreting forward interest rates: Sweden 1992-1994. NBER Working Paper 4871, National Bureau of Economic Research, 1994. [p1]