WHITE PAPER

Pricing the Term Structure of Interest Rates Using Linear Regressions in MATLAB

MathWorks®

Abstract

This white paper presents a detailed methodology for estimating and analyzing the term structure of interest rates using a regression-based framework. Drawing from the work of Adrian, Crump, and Moench (2013), the approach enables the estimation of affine term structure models without requiring numerical optimization. Instead, the model is estimated entirely through a sequence of linear regressions, making it computationally tractable even with high-dimensional yield data and multiple pricing factors. The methodology is implemented in MATLAB® using a function that processes historical yield data, performs principal component analysis, estimates factor dynamics, and decomposes observed yields into expectations and term premia. The practical implementation is illustrated using both U.S. Treasury yields and U.K. government bond data, demonstrating the model's flexibility and robustness.

Introduction

The term structure of interest rates encapsulates the relationship between the yield of zero-coupon bonds and their time to maturity. Modeling this relationship is central to understanding monetary policy expectations, macroeconomic risks, and fixed-income asset pricing. Traditional term structure models, particularly those grounded in affine frameworks, provide a theoretically consistent way to model interest rate dynamics and to compute term premia—the excess returns investors demand for holding long-term bonds instead of rolling over short-term instruments.

Affine models typically assume that bond yields are linear functions of underlying latent or observable state variables. The evolution of these state variables follows a vector autoregressive model, while the market prices of risk associated with each factor determine the term premia. In maximum likelihood implementations, such models are estimated by fitting observed yields directly, subject to the no-arbitrage conditions derived from dynamic asset pricing theory. However, these implementations are often computationally intensive and sensitive to specification assumptions.

Adrian, Crump, and Moench (2013) propose an alternative estimation methodology that sidesteps many of these difficulties. Their approach decomposes the model into three sequential linear regression steps. By focusing on excess bond returns and their relation to observable pricing factors, their method achieves empirical tractability and theoretical consistency. This paper implements their methodology in MATLAB and applies it to real-world yield data to demonstrate its empirical performance and interpretive power.

Theoretical Framework

Overview of the ACM Model

The ACM model assumes that bond yields are affine functions of a vector of state variables X_t, which are typically extracted from the yield curve using principal component analysis (PCA). These state variables evolve over time according to a first-order vector autoregressive process:

$$X_{t+1} = \mu + \Phi X_t + v_{t+1}$$

where μ is a vector of intercepts, Φ is a transition matrix, and v_{t+1} is a vector of shocks that are conditionally Gaussian with covariance matrix Σ.

The key insight of the ACM methodology is that the excess holding period return on a bond of maturity n is driven by its exposure to these shocks and by the market prices of risk. The excess return is modeled as:

$$r\,x^{(n-1)}{}_{t+1} = \beta^{(n-1)'}(\lambda_0 + \lambda_1 X_t) - \frac{1}{2}\left(\beta^{(n-1)'} \sum \beta^{(n-1)} + \sigma^2\right) + \beta^{(n-1)'} v_{t+1} + e_{t+1}{}^{(n-1)}$$

where $\beta^{(n-1)}$ is the factor loading for bond maturity n, λ_0 and λ_1 govern the market prices of risk, and $e_{t+1}^{(n-1)}$ is a return pricing error term.

This formulation implies that the term premium—the difference between the model-implied yield and the risk-neutral expected short rate—can be extracted from observed returns and factor innovations using regression techniques.

Three-Step Estimation Procedure

The ACM estimation proceeds in three steps:

1. **Estimation of Factor Dynamics**: A vector autoregression (VAR) is fitted to the pricing factors X_t. The residuals from this VAR provide estimates of factor innovations $\widehat{v}_{t+1}$.
2. **Excess Return Regression**: Excess returns on bonds of various maturities are regressed on lagged factors and contemporaneous factor innovations. This regression yields the factor loadings $\beta^{(n)}$, intercepts, and pricing errors.
3. **Estimation of Market Prices of Risk**: The estimated exposures to shocks and factor loadings are used in a cross-sectional regression to estimate the market prices of risk λ_0 and λ_1. These parameters define the stochastic discount factor used to price bonds and compute fitted yields and term premia.

This procedure yields closed-form estimators and allows for inference using asymptotic distributions derived from OLS theory. No numerical maximization is required, which contrasts with traditional affine term structure models estimated by likelihood-based methods.

MATLAB Implementation

We created a MATLAB implementation of the methodology using two datasets to illustrate the application of the ACM model. The first is a U.S. Treasury dataset based on Gürkaynak, Sack, and Wright (2007), while the second is a Bank of England dataset of zero-coupon yields.

The script loads the historical yield data, selecting a start date for analysis (e.g., January 1985), and interpolates missing values. The resulting `yields` variable is a timetable containing yield observations across a grid of maturities. The function then carries out the estimation of the model. It performs PCA on the yields to extract factors, estimates the VAR for factor dynamics, and regresses excess returns to infer loadings and market prices of risk.

The output `decomposition` includes the fitted yield curve (`Fitted`), the risk-neutral expected yields (`RiskNeutralExpected`), and the estimated term premium (`TermPremium`) for each maturity.

The decomposition results can be visualized across different maturities to assess model fit and economic interpretation. For the 2-, 5-, and 10-year maturities, one can plot the actual yield, fitted yield, risk-neutral yield, and term premium over time:

This allows the analyst to see how well the model fits yields and to study the time variation in the term premium – a key input for understanding risk compensation in bond markets.

Application to U.K. Yield Data

The same model is applied to zero-coupon yields from the U.K. government bond market. The zero-coupon data is imported from a spreadsheet and preprocessed in a similar way as the U.S. data.

The same estimation and visualization procedure is used, demonstrating the model's portability across different markets.

A time-series chart showing the evolution of the term premium for the 2-, 5-, and 10-year U.K. government bonds is shown below. This figure can reveal episodes of monetary policy tightening or loosening as reflected in the risk compensation required by investors.

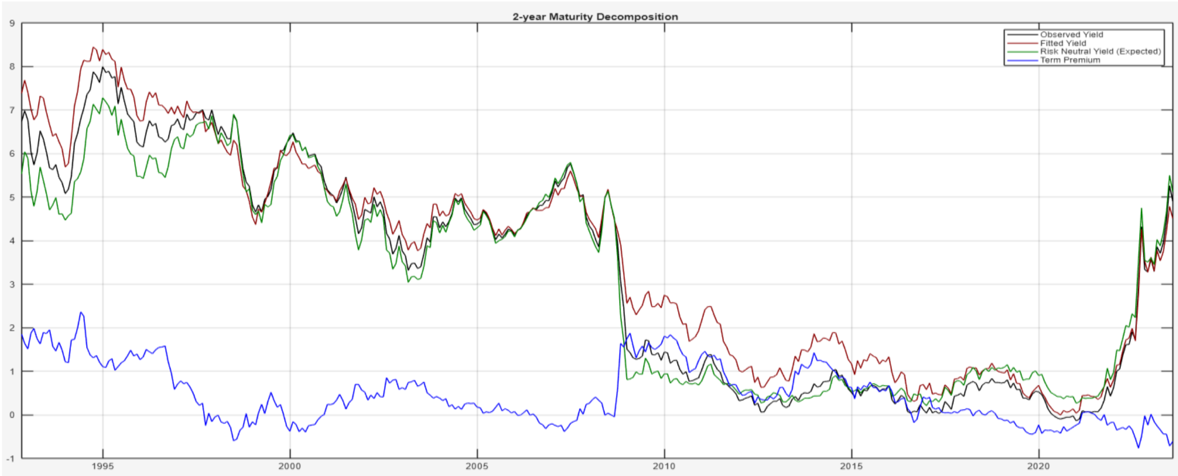
2-year Maturity Decomposition
Observed Yield
Fitted Yield
Risk Neutral Yield (Expected)
Term Premium
9
8
7
6
5
4
3
2
1
0
-1
1995
2000
2005
2010
2015
2020

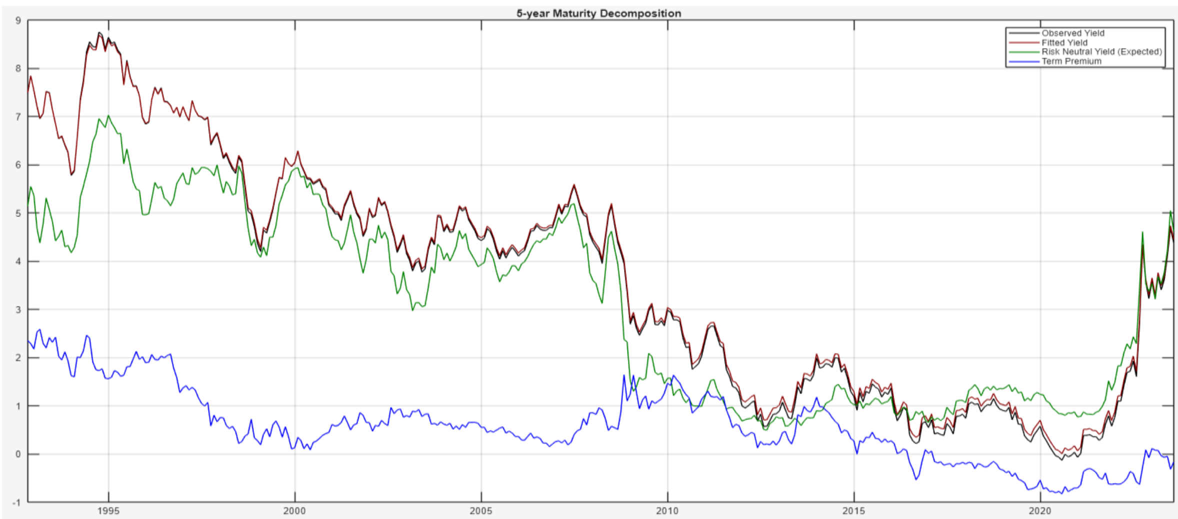
5-year Maturity Decomposition
Observed Yield
Fitted Yield
Risk Neutral Yield (Expected)
Term Premium
9
8
7
6
5
4
3
2
1
0
-1
1995
2000
2005
2010
2015
2020

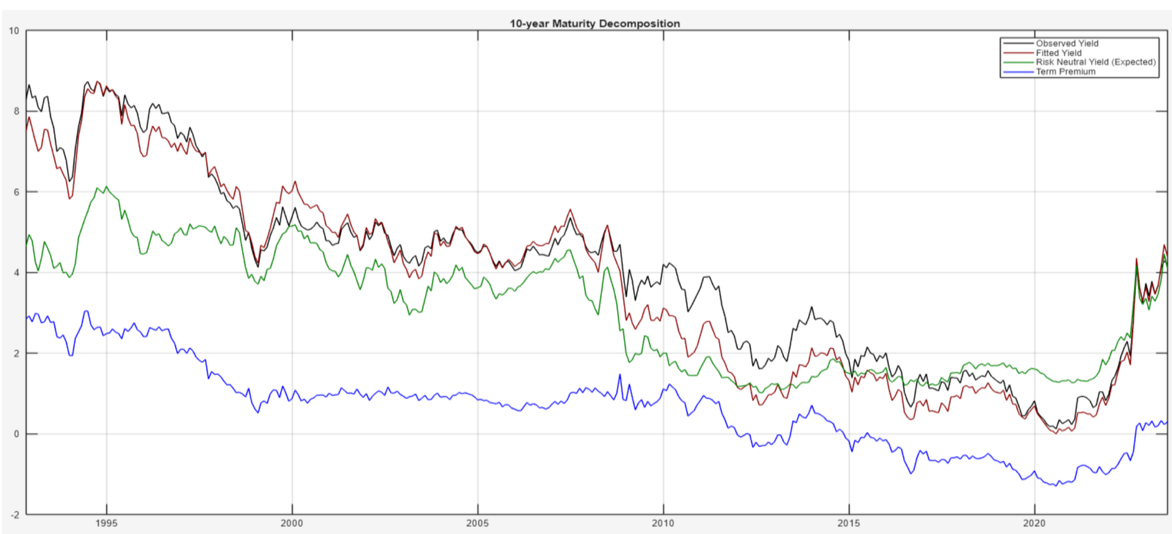
10-year Maturity Decomposition
Observed Yield
Fitted Yield
Risk Neutral Yield (Expected)
Term Premium
10
8
6
4
2
0
-2
1995
2000
2005
2010
2015
2020

The root mean squared error (RMSE) between the observed and fitted yields provides a summary metric of model accuracy. This metric can be plotted as a function of maturity to determine whether the model fits certain parts of the yield curve better than others. Generally, affine models tend to fit short and intermediate maturities more closely than the long end, where pricing errors and idiosyncratic factors may dominate.

Statistical Inference and Model Diagnostics

Estimation Accuracy and Standard Errors

A core strength of the ACM methodology lies in the transparency and reliability of inference. All parameter estimates—including factor loadings $\beta^{(n)}$, market price of risk parameters λ_0 and λ_1, and the VAR dynamics of pricing factors—are obtained through ordinary least squares regressions. As a result, standard errors are analytically tractable and asymptotically valid under standard regularity conditions.

The asymptotic distribution of the estimated market price of risk vector $\widehat{\Lambda} = [\widehat{\lambda}_0, \widehat{\lambda}_1]$ is given by:

$$\sqrt{T}\,(\widehat{\Lambda} - \Lambda) \xrightarrow{d} \mathcal{N}(0, V_\Lambda)$$

The matrix V_Λ can be estimated consistently using the regression residuals and factor innovations. This allows for:

- Hypothesis testing on whether specific factors are priced.
- Wald tests to assess joint significance of parameters.
- Confidence intervals for model-implied term premia.

For example, one might test whether slope risk is priced by testing the null hypothesis $H_0 : \lambda_{\text{slope}} = 0$ using the Wald statistic: $W = \widehat{\lambda}' V_{\widehat{\lambda}}^{-1} \widehat{\lambda} \sim \chi_k^2$

where k is the number of restrictions (e.g., the number of rows of λ_1 corresponding to slope).

Goodness of Fit: RMSE and Yield Pricing Errors

Beyond coefficient significance, the model's in-sample accuracy is evaluated using the root mean squared error (RMSE) between observed and fitted yields. In practice, the ACM model achieves RMSEs on the order of 1–2 basis points for intermediate maturities. Standard deviations of yield pricing errors remain well below 5 basis points, indicating a tight fit to observed term structures.

Autocorrelation in pricing errors is another diagnostic. The ACM framework shows that while yield pricing errors exhibit serial correlation—particularly at longer maturities—return pricing errors do not, affirming the model's validity for return-based estimation.

Macroeconomic Interpretation of Term Premia

Components of Long-Term Yields

The decomposition provided by the ACM model distinguishes between the **expected path of future short rates** (the risk-neutral yield), which reflects investor expectations of monetary policy and the **term premium**, which captures the additional compensation required by investors for interest rate and macroeconomic uncertainty.

This separation is particularly useful for interpreting yield movements around policy shifts. For instance, a decline in long-term yields during a period of rising short rates may reflect a falling term premium rather than an expectation of future rate cuts.

Time Variation in Risk Compensation

By examining the estimated λ_1 matrix, one can observe how the price of risk varies over time and in response to macroeconomic variables. For example, Adrian et al. (2013) find that the second and fifth principal components—often associated with slope and a "macro risk" factor—play significant roles in driving expected excess returns. A steep yield curve or elevated macro uncertainty increases the compensation investors demand for bearing duration risk.

In MATLAB, these interpretations can be made concrete by analyzing the time series of the term premium and correlating it with external indicators such as inflation expectations, real GDP growth and central bank balance sheet expansions.

Such overlays help assess whether changes in the term premium are aligned with macroeconomic fundamentals or driven by technical factors.

Comparative Analysis with Alternative Models

Likelihood-Based Affine Models

Traditional affine term structure models are estimated using maximum likelihood techniques, often requiring complex numerical optimization and Kalman filtering to handle latent factors and measurement error. While theoretically elegant, these methods suffer from high computational costs, sensitivity to initial conditions and convergence criteria, difficulty scaling to more than three or four pricing factors.

In contrast, the ACM model estimates all parameters using linear regressions and observable factors. This approach preserves economic relationships and gains in scalability and diagnostic transparency.

In benchmark tests, ACM estimation achieves comparable in-sample fit to ML models while being orders of magnitude faster to estimate. Moreover, it is less susceptible to overfitting due to fewer moving parts and a more modular structure.

Cochrane-Piazzesi Return Forecasting Model

The Cochrane-Piazzesi (CP) model constructs a single linear combination of forward rates that forecasts bond excess returns. This return-forecasting factor is highly interpretable and parsimonious but assumes a very specific structure for the price of risk: Level is the only priced factor, and time variation in risk premia is entirely driven by a single predictor.

When re-estimated within the ACM framework, the CP factor is found to be statistically significant, particularly in explaining the time variation of level risk. However, the full five-factor ACM model outperforms the CP specification in cross-sectional pricing accuracy, out-of-sample forecasting of future short rates and fitting long-maturity yields beyond the calibration window.

Conclusion

The regression-based affine term structure modeling approach introduced by Adrian, Crump, and Moench provides a powerful, scalable, and interpretable framework for understanding the yield curve and its components. When implemented in MATLAB, the model yields tight fits, plausible term premium dynamics, and economically intuitive results, all while avoiding the numerical complexity of traditional ML-based approaches.

Its ability to handle observable factors, macro-financial predictors, and excess return dynamics makes it a valuable tool for central banks, institutional investors, and academics interested in monetary policy transmission, risk pricing, and yield forecasting.

The MATLAB ecosystem—featuring advanced econometric functionality, effective visualizations, and integration with databases— makes it ideal for deploying this framework in production environments or research workflows.

Next Step

References

Adrian, T., Crump, R. K., & Moench, E. (2013). Pricing the term structure with linear regressions. *Journal of Financial Economics*, 110(1), 110–138.

Cochrane, J. H., & Piazzesi, M. (2005). Bond risk premia. *American Economic Review*, 95(1), 138–160.

Gürkaynak, R. S., Sack, B., & Wright, J. H. (2007). The US Treasury yield curve: 1961 to the present. *Journal of Monetary Economics*, 54(8), 2291–2304.