Impact of covariance matrix estimator choice on yield curve decomposition with PCA

Grant Holtes

November 26, 2023

Summary

This investigation evaluates the impact of different covariance matrix estimation methods when used as an input for principle component analysis in the context of yield curve decomposition, focusing on the economic interpretations of Level, Slope, and Curvature components. It was found that the constant correlation shrinkage estimator consistently outperforms others, including the sample covariance matrix. The investigation also reveals that short time periods yield less reliable results, with the issue attributed to the length of time sampled rather than sample sizes.

1 Introduction

The decomposition of yield curves using principle component analysis (PCA) is a common procedure, allowing for yield curve dynamics to be explained in a small number of underlying drivers. The first three principle components are commonly used and are given economic interpretations as in Litterman and Scheinkman, 1991 and Kavir Patel, 2018.

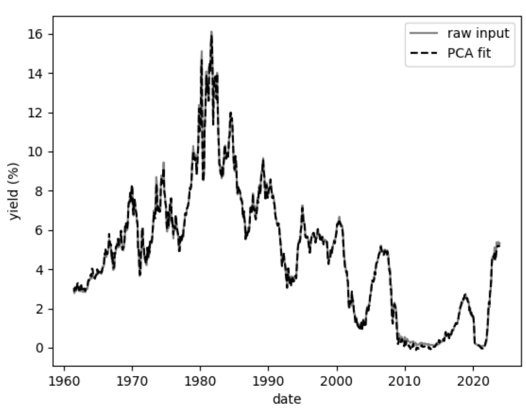

Figure 1: US zero coupon bond yield curve reconstructed with 3 component PCA, 12 month time to maturity series shown

PC1: Level: The first principle component tends to have equal values across all maturities, so controls the overall level of the curve.

PC2: Slope: The second component usually has a positive relationship with maturity, which allows this component to determine the slope of the curve and the bulk of the differences in yield between short and long maturities.

PC3: Curvature: The third and usually last component tends to determine how bowed the curve is, leading to it being referred to as curvature.

These three components can be visualised by plotting the loadings on each principle component for each maturity, the eigenvectors of the covariance matrix, against maturity, as in Figure 2.

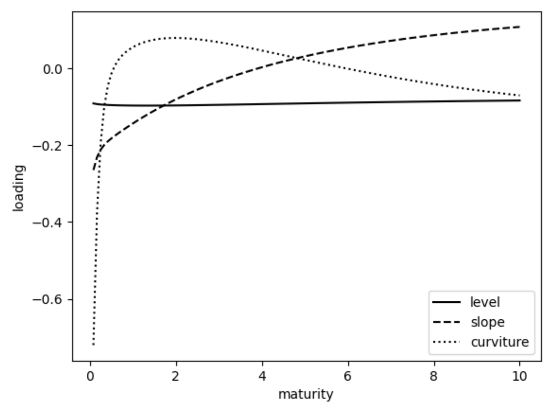

Figure 2: Principle component loadings on US zero coupon bond time data, entire data history

The use of PCA to decompose the yield curve into factors with distinct economic interpretations raises some questions on how the reliability of the decomposition, as the decomposition is completely informed by historical data.

This question is explored by Rodger Lord, 2007, where the authors investigate the circumstances under which the economic interpretations are consistent with their Level, Slope and Curvature designations, concluding that there are some years in which there are no economic interpretations for the components. While Rodger Lord, 2007 use the shape and sign changes in the eigenvectors to determine if the economic interpretation is valid, Hou, 2022 uses a comparison against the model developed by Nelson and Siegel, 1987, which explicitly models the components, to test if the PCA based components are poorly estimated, also concluding that there are time periods where the interpretations of the components are inconsistent with their usual interpretations.

This investigation aims to determine impact of the choice of covariance estimator on the estimation of the principle components of bond yields and whether specific estimators provide more consistency on historic data.

2 Approach

Two approaches are explored to quantify how well estimated the principle components are:

Sign change of eigenvectors: This borrows the method and sign change rules used by Rodger Lord, 2007, which checks if "first three factors have respectively zero, one and two sign changes", in which case the level, slope and curvature are assumed to be estimated in accordance with their economic interpretations. Applying these rules to the results in Figure 2 we can see that the decomposition is valid

Visual Inspection: Inspection of graphs of eigenvectors against maturity to try find any insights that may have been missed by either of the other methods.

These approaches are applied to the decomposition of a set of rolling sub-ranges of yields from the US treasury yields data from the Board of Governors of the Federal Reserve System (Reserve, n.d.).

3 Results

3.1 Eigenvector sign change analysis

The choice of covariance estimator had a small but consistent impact on the reliability of PCA in generating eigenvectors that are consistent with the sign change assumptions. This can be seen in Table 1, where the constant correlation shrinkage covariance estimator consistently outperformed the sample covariance estimator while the simple shrinkage estimator consistently underperformed.

The Ledoit-Wolf estimator performed identically to the sample covariance estimator, suggesting that the optimiser had (correctly) inferred that there was little to no benefit in shrinkage towards the diagonal matrix, as in the simple shrinkage estimator.

months	covariance estimator	proportion valid
12	constant correlation shrinkage estimator	0.74
12	ledoit wolf shrinkage estimator	0.626
12	sample covariance matrix	0.626
12	simple shrinkage estimator	0.512
24	constant correlation shrinkage estimator	0.893
24	ledoit wolf shrinkage estimator	0.802
24	sample covariance matrix	0.802
24	simple shrinkage estimator	0.678
48	constant correlation shrinkage estimator	0.932
48	ledoit wolf shrinkage estimator	0.897
48	sample covariance matrix	0.897
48	simple shrinkage estimator	0.795
96	constant correlation shrinkage estimator	1.0
96	ledoit wolf shrinkage estimator	0.982
96	sample covariance matrix	0.982
96	simple shrinkage estimator	0.917

Table 1: Proportion of test periods with consistent sign change patterns in the PCA eigenvectors, by covariance estimator and test period length

The other, more intuitive pattern in Table 1 is that performance across all estimators improved with longer sample windows, illustrating how the decomposition is unreliable over shorter time frames. The patterns in the performance of the estimators can also be explored over time, by plotting whether or not each estimation is valid on a rolling basis through the dataset, as in Figure 3 and Figure 4. These plots how how the errors in the estimation are highly correlated across estimators and across sample size ranges, with most of the estimators performing poorly on the same time periods.

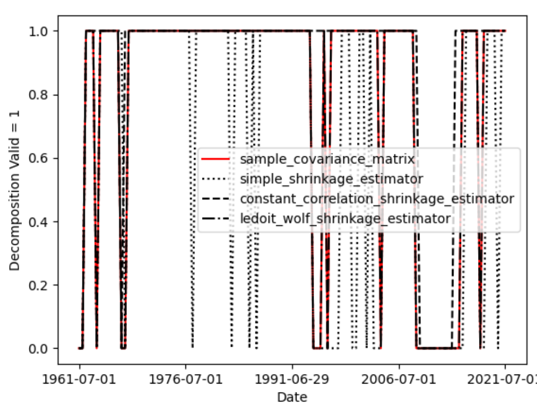

Figure 3: Eigenvector consistency with economic interpretations over time, 24 month rolling window

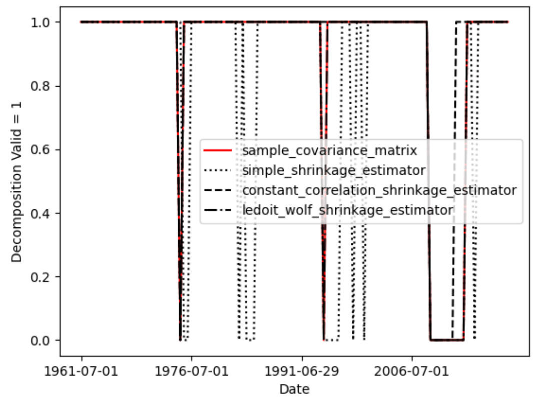

Figure 4: Eigenvector consistency with economic interpretations over time, 48 month rolling window

These results raise a question: Is the poor performance on small sample sizes due to the short time period sampled, or the small number of samples? This can be investigated by re-running the analysis with the same sample sizes, but with quarterly rather than monthly data, so each slice contains the same number of observations but three times the elapsed time. Comparing these quarterly results in Table 2 to the monthly results in Table 1 suggests that the poor performance on short time windows is largely due to the length of time rather than the number of samples, likely due to insufficient variation in yields over the shorter periods of one-to-two years to provide the required information.

quarters	covariance estimator	proportion valid
12	constant correlation shrinkage estimator	0.85
12	ledoit wolf shrinkage estimator	0.825
12	sample covariance matrix	0.825
12	simple shrinkage estimator	0.65
24	constant correlation shrinkage estimator	0.947
24	ledoit wolf shrinkage estimator	0.895
24	sample covariance matrix	0.895
24	simple shrinkage estimator	0.842

Table 2: Proportion of test periods with consistent sign change patterns in the PCA eigenvectors, by covariance estimator and test period length

3.2 Visual Inspection

We can use visual inspection to understand the results of the sign change analysis in more detail and to find other inconsistencies in the PCA results that were not found by the sign change measure.

Figure 5 and Figure 6 show examples of eigenvectors in one of the sample periods that the constant correlation shrinkage estimator produced valid sign changes but the sample covariance matrix did not. We can see that the cause of the issue is the 3 sign changes in the curvature vector in the sample covariance graph.

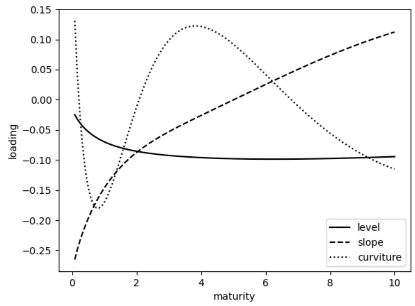

Figure 5: Eigenvectors for July 1997 to April 2000, sample covariance estimator, invalid sign changes

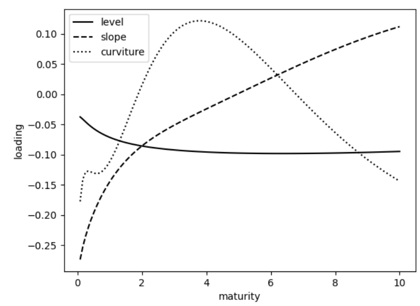

Figure 6: Eigenvectors for July 1997 to April 2000, constant correlation shrinkage covariance, valid sign changes

We can also see some interesting ways that the eigenvectors can be correct from a sign change perspective, but lack the properties to make them useful in practice. The sample covariance matrix tended to produce very flat eigenvectors when estimated on short time spans of data, as in Figure 7. This lack of structural difference between the eigenvectors after the first few months of maturity limits their usefulness, as they suggest near equal impacts from the principle components on medium and long term bonds.

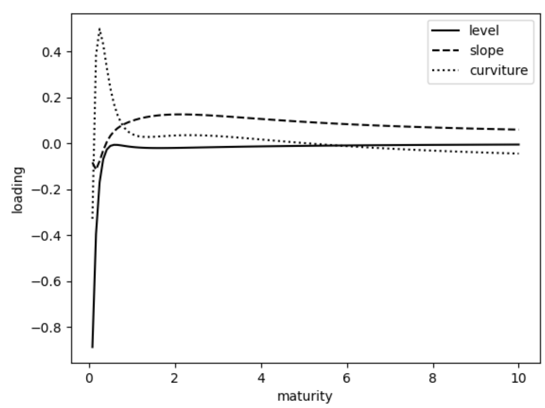

Figure 7: Eigenvectors for August 1963 to August 1964, , sample covariance estimator. Sign changes suggest a correctly estimated set of eigenvectors, but produce very flat and unhelpful curves.

4 Conclusion

The main findings are:

Decompositions are unreliable over short periods: Short time periods do not provide the required information to estimate the principle component loadings reliably.

Time period matters more than samples: Poor results over short time periods seem attributable to a lack of time, rather than the small sample sizes. The same small sample sizes spread over a longer period performed much better, and practically as well as more frequent samples over the same longer period.

The constant correlation shrinkage estimator performs well: The constant correlation shrinkage estimator consistently outperformed the other estimators, including the sample covariance matrix.

The simple shrinkage estimator performs poorly: The simple shrinkage estimator consistently underperformed the sample covariance matrix. This lead to other estimators such as the Ledoit-Wolf estimator, which also shrinks between the sample and diagonal matrix, to preform poorly as well, with limited shrinkage benefits.

References

Hou, K. (2022). *Inference for model misspecification in interest rate term structure using functional principal component analysis.*

Kavir Patel, G. W. v. V., Ashfaaq Mohamed. (2018). A regression and comparative study of united states and south african yield curves using principal component analysis. *South African Journal of Economic and Management Sciences*, 1–15.

Litterman, R., & Scheinkman, J. (1991). Common factors affecting bond returns. *Journal of fixed income.*

Nelson, C. R., & Siegel, A. F. (1987). Parsimonious modeling of yield curves. *Journal of business*, 473–489.

Reserve, F. (n.d.). *Yield curve models and data* [Data in CSV format].

Rodger Lord, A. P. (2007). Level–slope–curvature – fact or artefact? *Applied Mathematical Finance.*