



UNIVERSITÀ POLITECNICA DELLE MARCHE
Dipartimento di Scienze Economiche e Sociali

**A ROTATED DYNAMIC FACTOR MODEL FOR THE
YIELD CURVE: SQUEEZING OUT INFORMATION WHEN
IT MATTERS**

Chiara Casoli Riccardo (Jack) Lucchetti

QUADERNI DI RICERCA n. 503*
ISSN: 2279-9575

January 2026

Comitato scientifico:

Giulia Bettin
Marco Gallegati
Stefano Staffolani
Alessandro Sterlacchini

Collana curata da Alberto Russo

Abstract

The yield curve is widely regarded as a powerful descriptor of the economy and market expectations. A common approach to its statistical representation relies on a small number of factors summarizing the curve, which can then be used to forecast real economic activity.

We argue that optimal factor extraction is crucial for retrieving information when considering an *approximate* factor model. By introducing a rotation of the model including cointegration, we reduce cross-sectional dependence in the idiosyncratic components.

This leads to improved forecasts of key macroeconomic variables during periods of economic and financial instability, both in the US and the euro area.

JEL Class.: C32, C53, E43

Keywords: Yield curve, Nelson-Siegel model, Dynamic Factor Model, cointegration, forecasting

Indirizzo: Chiara Casoli (**corresponding author**): InsIDE Lab, DiEco, Università degli Studi dell'Insubria, Fondazione Eni Enrico Mattei
Riccardo (Jack) Lucchetti: DiSES, Università Politecnica delle Marche

A rotated Dynamic Factor Model for the yield curve: squeezing out information when it matters*

Chiara Casoli Riccardo (Jack) Lucchetti

1 Introduction

Two fundamental facts about the yield curve have been known for a long time: *i*) it embodies information on the state of the economy and market expectations that can be very useful for several purposes, notably forecasting or identification of monetary policy shocks; *ii*) its structure can be described, at a given point in time, by a limited number of factors.

The profound interactions of macroeconomic factors with the yield curve has led many to monitor the term structure of interest rates for predicting recessions and creating better proxies for expectations of future economic conditions (see Chinn and Ferrara, 2024). This builds on a long tradition linking the slope of the yield curve to future economic conditions (Estrella and Hardouvelis, 1991).

In general, as explained in Coroneo et al. (2016), the short end of the yield curve is more closely connected with the policy instruments of the central banks, whereas the average level of the yield curve usually co-moves with broader macroeconomic forces, such as the inflation rate. Finally, the spread of long versus short rates is associated with business cycle conditions. These interactions are often analysed in order to predict economic recessions (Bordo and Haubrich, 2024; Minesso et al., 2022), interest rates (Caruso and Coroneo, 2023), financial crises (Bluwstein et al., 2023) and the business cycle (Han et al., 2021).

*We wish to thank Gianni Amisano, Matteo Barigozzi, Andrea Bellucci, Emanuele Ciola, Laura Coroneo, Gialluca Cubadda, Matteo Luciani, Barbara Rossi, Esther Ruiz, Evi Salamaliki, Rafaello Seri, Yiannis Venetis and Mark Watson. We are also indebted to all the participants of the 6th Vienna Workshop on High-Dimensional Times Series in Macroeconomics and Finance, the 11st Italian Congress of Econometrics and Empirical Economics, the 3rd Vienna Workshop on Economic Forecasting and the 2025 International Association for Applied Econometrics Conference in Turin. Chiara Casoli gratefully acknowledges financial support from the "Department of Excellence 2023–2027" provided by the Italian Ministry of University and Research.

Modelling the yield curve is challenging because of its complexity and the dynamic nature of the factors influencing it, which can also be unobservable. Importantly, different modelling choices can lead to heterogeneous forecasting performances (see, e.g., Caldeira et al., 2025).

One of the most commonly accepted way of modelling the yield curve relies on the fact that the entire structure can be described by a small number of factors. Examples are the famous Nelson and Siegel (1987) model, which synthesizes all the information contained in several yields in only three factors governing the shape of the curve, and the extension to four factors proposed by Svensson (1994).¹ Diebold and Li (2006); Diebold et al. (2006) extend the Nelson-Siegel model by considering *dynamic* factors, opening the way to dynamic factor models (DFM) as the preferred estimation tool for modelling the yield curve.

In this paper, we revisit the important contributions by Diebold and Li (2006) and Diebold et al. (2006) in the light of more recent developments in the literature on Dynamic Factor Models and provide evidence on how differences in the methods for extracting yield curve factors are reflected in the estimated components, and thereafter, in forecasting performance. Specifically, Diebold and Li (2006) express yields as combinations of three common dynamic factors governing the level, slope and curvature of the curve. In this framework, the loading matrix is constrained and depends non-linearly on a scalar unknown parameter, universally denoted as λ .²

One of the main ideas in the present contribution is that the cointegration properties of the observed yields can be used to perform a rotation of the observable terms: we show that taking the cointegration structure of the yields into account leads to a significant reduction of the cross-sectional correlation of the idiosyncratic components of the DFM.³ The advantages of embedding cointegration in a DFM are highlighted in Barigozzi and Luciani (2019); Casoli and Lucchetti (2022). Our setup considers a similar background to the Casoli and Lucchetti (2022) work, in which cointegration is assumed among the observable variables rather than the factors.

The importance of reducing the cross-correlation of the idiosyncratic components relies on the fact that more information is included in the

¹Recently, Wahlstrøm et al. (2022) have demonstrated that the Nelson-Siegel parameter estimates are more stable with respect to those of the Svensson alternative.

²Notably, the literature on DFMs with constrained loading matrices is evolving rapidly and is closely related to this setup. Examples include multilevel DFMs (Breitung and Eickmeier, 2015; Choi et al., 2018) or matrix/tensor DFMs (see Chen et al., 2024; Yu et al., 2024, as examples)

³Note that this is an approximate DFM, so that the variance-covariance matrix of the idiosyncratic term is not diagonal.

common component, therefore leading to a better estimation of the factors space. Fresoli et al. (2023) demonstrate that ignoring cross-correlated idiosyncratic components has implications in terms of forecasting, as it can underestimate prediction intervals.

Additionally, we compare the forecasting performance of two sets of yield-curve factors extracted from an original (unrotated) model and from a rotated model that incorporates cointegration. Using these factors, we forecast several macroeconomic variables for both the United States and the Euro area and assess predictive performance by comparing predictive log-likelihoods.

We find that the two models exhibit very similar forecasting ability outside periods of economic instability and financial stress. However, during episodes of large turbulence, such as the 2008 global financial crisis and the Covid-19-related recession, the rotated model consistently outperforms the classic specification. This evidence suggests that incorporating cointegration into yield-curve factor models is particularly relevant for predicting recessions.

Finally, as a secondary result, we also estimate by Maximum Likelihood the parameter λ governing the Nelson-Siegel loadings. The literature has typically relied on pre-set values, but our estimates indicate much smaller values of λ both for the US and the Euro area datasets. This finding is consistent with the persistently flatter yield curves observed during the zero lower bound (ZLB) period. While this is not the core focus of the paper, it provides new empirical evidence that may inform future work on the evolution of the term structure under unconventional monetary conditions.

The structure of the paper is as follows: Section 2 introduces the motivation and research questions, Section 3 explains the way cointegration can be embedded in a DFM, Section 4 describes our dataset, while our results are in Sections 5 and 6. Section 7 concludes.

2 Motivation and preliminary evidence

As we anticipated in the previous section, the idea that the yield curve contains information that may be useful for predicting macro variable has been used in many cases for a long time. The possibility of using such information is even more tempting if the entire structure of the yield curve can be summarised by a few factors, the Nelson-Siegel model being the most popular choice.

However, it can be conjectured that the predicting power of the yield curve factors may not be uniform through time: several practitioners have

recently been concerned on the apparently decreasing relevance of the yield curve as a predictor for recessions (see for instance Chinn and Ferrara, 2024). A possible explanation relates the flattening of the curve, with also a decrease in the yields volatility, to the ZLB (Opschoor and van der Wel, 2024). A related and growing literature documents how conventional models often lose explanatory power in ZLB environments (Rossi, 2021; Wu and Xia, 2016). Therefore, extracting the relevant information from the yield curve is a crucial task for several purposes, and the challenge concerning whether there is a decreasing predictive power of the yield curve is open and difficult. That said, it is quite natural to think that, if the added value of such information is linked to its ability to capture agents' expectations, its importance should increase in times of economic turmoil.

In order to check for this possibility, we ran a very simple preliminary experiment, using 30 selected maturities from the US yield curve and computing the first three principal components.⁴ For this analysis, we rely on the standard Nelson-Siegel model with a dynamic factor structure (For the notation, we refer the reader to Section 3). The results conform very well with economic intuition: the first three principal components contain 99.96% of the total information and the structure of the loading matrix is strikingly close to what one would expect for the “level”, “slope” and “curvature” factors in the Nelson-Siegel model, as shown in Table 1.⁵ It would seem that the Nelson-Siegel factors may in fact provide an excellent summary of the information contained in the yield curve.

Next, we used the estimated factors in a forecasting model like

$$y_t = \mu + \sum_{i=1}^{12} \alpha_i y_{t-1} + \sum_{i=1}^{12} \beta'_i \mathbf{F}_{t-i} + \varepsilon_t \quad (1)$$

where y_t is a macroeconomic variable and \mathbf{F}_t are the three principal component extracted above. Model (1) was estimated on a rolling window of 240 months and the p -value for the Granger-causality hypothesis $\beta_1 = \beta_2 = \dots = \beta_{12} = 0$ was computed, mainly as a descriptive statistic of the predictive power of \mathbf{F}_t for y_t , where of course smaller values imply greater predictive power.

Results for a few US macroeconomic series (taken from FRED), are shown in Figure 1. It is quite evident that the predictive power of principal components is very high during the main episodes of economic instability:

⁴The data are available from Jing Cynthia Wu's website, as reported in Section 4.

⁵In principle, even if the Nelson-Siegel model was *exactly* true, principal components would just be a basis for the space spanned from the three factors. This distinction, however, is immaterial for the results in this section.

	PC1	PC2	PC3
3 months	0.176	0.290	0.458
6 months	0.178	0.280	0.341
9 months	0.179	0.266	0.238
12 months	0.180	0.250	0.145
15 months	0.181	0.230	0.063
18 months	0.182	0.208	-0.004
21 months	0.183	0.187	-0.056
24 months	0.183	0.166	-0.099
27 months	0.184	0.146	-0.134
30 months	0.184	0.126	-0.161
33 months	0.185	0.108	-0.179
36 months	0.185	0.091	-0.192
42 months	0.186	0.058	-0.209
48 months	0.186	0.028	-0.213
54 months	0.186	-0.001	-0.205
60 months	0.186	-0.026	-0.189
66 months	0.186	-0.049	-0.173
72 months	0.185	-0.071	-0.150
78 months	0.185	-0.091	-0.123
84 months	0.185	-0.111	-0.093
90 months	0.184	-0.130	-0.062
96 months	0.184	-0.145	-0.033
102 months	0.184	-0.160	-0.006
108 months	0.183	-0.173	0.022
120 months	0.182	-0.201	0.085
132 months	0.181	-0.225	0.138
144 months	0.180	-0.244	0.179
156 months	0.179	-0.257	0.211
168 months	0.179	-0.267	0.235
180 months	0.178	-0.273	0.246

Table 1: PC loadings on US yields, 30 maturities.

the 2008 financial crisis, the COVID outbreak and the Russia-Ukraine conflict. This effect is not uniform across all the macro series, but the fact that for all macro series the predictive power of PCs was rather limited during the 2012–2018 period is apparent.

These observations give rise to two fundamental questions, which constitute the central research questions addressed in this article.

1. Is there an optimal way to condensate the information contained in the yield curve so as to make it useful for macroeconomic forecasting?
2. Is it possible to devise a metric for the effectiveness of these factors as predictors that takes into account its possibly time-varying nature?

The first question will be explored next, in Section 3, while our approach to the second one will be described in subsection 5.3.

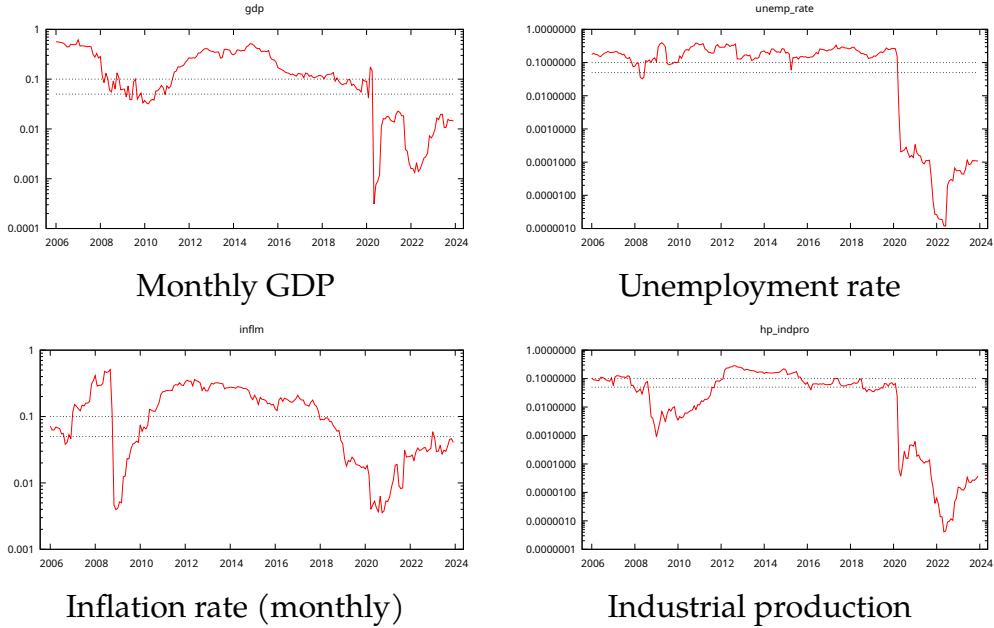


Figure 1: Predictive power of principal components

3 The rotated Nelson-Siegel model

3.1 The classic Nelson-Siegel model

The Nelson-Siegel Nelson and Siegel (1987) model describes the yield curve for a set of risk-free bonds as a function of their maturity. In the Diebold-Li fashion (Diebold and Li (2006); Diebold et al. (2006)), the yield of a bond with maturity τ is expressed as

$$y(\tau) = \beta_0 l(\tau) + \beta_1 s(\tau) + \beta_2 c(\tau), \quad (2)$$

where the three components are known as “level”, “slope” and “curvature”. Each is a nonlinear function of τ as follows:

- $\beta_0 = 1$: constant across the maturity spectrum, models parallel yield curve shifts (long-term factor).
- $\beta_1 = \frac{1-e^{-\lambda\tau}}{\lambda\tau}$: loading starts at 1 but decays to 0 with maturity. Interpreted as the (negative of the) slope of yield curve (short-term factor).
- $\beta_2 = \frac{1-e^{-\lambda\tau}}{\lambda\tau} - e^{-\lambda\tau}$: loading starts at 0, increases, and then decays to zero. It gives maximum weight to intermediate maturities (medium-term factor).

The representation of yields via a Dynamic Factor Model (DFM) is:

$$\mathbf{y}_t = \Lambda(\lambda) f_t + e_t, \quad (3)$$

where f_t is a 3-element vector containing the three factors:

$$f'_t = [L_t \quad S_t \quad C_t],$$

and the matrix of loadings Λ , in which the only unknown parameter is λ , has a precise structure:

$$\Lambda = \begin{bmatrix} 1 & \frac{1-e^{-\lambda\tau_1}}{\lambda\tau_1} & \frac{1-e^{-\lambda\tau_1}}{\lambda\tau_1} - e^{-\lambda\tau_1} \\ 1 & \frac{1-e^{-\lambda\tau_2}}{\lambda\tau_2} & \frac{1-e^{-\lambda\tau_2}}{\lambda\tau_2} - e^{-\lambda\tau_2} \\ \vdots & & \\ 1 & \frac{1-e^{-\lambda\tau_n}}{\lambda\tau_n} & \frac{1-e^{-\lambda\tau_n}}{\lambda\tau_n} - e^{-\lambda\tau_n} \end{bmatrix}.$$

The fact that $\mathbf{y}_t \sim I(1)$ in most cases has been traditionally handled by differencing variables. However, as shown by Barigozzi and Luciani (2019); Casoli and Lucchetti (2022), unit roots can be accommodated if cointegration relationships are present.

3.2 The rotated model

As the (local) expectation hypothesis theory of the term structure of interest rates entails,⁶ we assume \mathbf{y}_t to be cointegrated, with one common trend and $n - 1$ spreads.⁷ Although the yields are $I(1)$, spreads are stationary by definition. This idea has been pursued in countless empirical applications, especially in the late 1990s, with an equally impressive number of variations on the theme. In this paper, we simply postulate the existence of a valid VECM representation for \mathbf{y}_t as

$$\Gamma(L)\Delta\mathbf{y}_t = \boldsymbol{\mu}_0 + \boldsymbol{\alpha}\boldsymbol{\beta}'\mathbf{y}_t + \varepsilon_t, \quad (4)$$

where the spreads are $\mathbf{s}_t = \boldsymbol{\beta}'\mathbf{y}_t$.

⁶See eg Brand and Cassola (2004), p. 819.

⁷To check if yields are in fact $I(1)$, we implemented a set of ADF tests for unit roots, which confirmed our hypothesis. However, there is not consensus on whether the interest rates should be considered as stationary or non-stationary processes. In a recent contribution, Rogoff et al. (2024) argue that they should be stationary over the long term. However, when persistence in finite sample data is very high, it may be a better choice to impose the unit root even if the DGP is in fact $I(0)$ with high persistence (see Di Iorio et al., 2016, for a similar discussion on near- $I(2)$ stochastic processes).

In principle, there are two main ways to represent β :

$$\beta_1 = \begin{bmatrix} 1 & 1 & \cdots & 1 \\ -1 & 0 & \cdots & 0 \\ 0 & -1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & -1 \end{bmatrix} \quad \beta_2 = \begin{bmatrix} 1 & 0 & \cdots & 0 \\ -1 & 1 & \cdots & 0 \\ 0 & -1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 \\ 0 & 0 & \cdots & -1 \end{bmatrix},$$

that is, all spreads are computed with respect to one yield, or the spreads are expressed as difference between bonds of similar maturity. In this paper, we use β_2 , which gives rise to what we call “adjacent” spreads. The other choice is possible, but was found to yield less satisfactory results.

Using representation (4), a “common trend” term m_t can be defined by combining the yields \mathbf{y}_t with a vector φ that does not belong to the space spanned by the cointegration matrix β :

$$m_t = \varphi' \mathbf{y}_t$$

There are two possible alternatives for φ , based on popular choices for trend-cycle decomposition in the cointegration literature.

- The Kasa (1992) decomposition: in this case, $\varphi = \beta'_\perp$. Under the stationarity assumption for the spreads, β'_\perp is a multiple of ι , and m_t is the simple average of all the rates.
- The Gonzalo and Granger (1995) decomposition, where $\varphi = \alpha'_\perp$. We have two choices:
 1. assume that one of the rates is weakly exogenous so that α_\perp is just a selection vector picking up that rate;
 2. estimate α via OLS and compute α_\perp from there; in this case, m_t is a weighted average of all the rates (possibly, with negative weights).

We use the Kasa decomposition, which produces more interpretable results and mitigates possible inferential problems when estimating α : together with the theory-based choice we make for β , this means that the long-run matrix $\Pi = \alpha\beta'$ in equation (4) contains no estimated elements.

Note that Equation (2) can be used to represent the spreads and the trend as

$$\mathbf{s}_t = \beta' \mathbf{y}_t = \beta' \Lambda f_t + v_t \quad (5)$$

$$m_t = \varphi' \mathbf{y}_t = \varphi' \Lambda f_t + u_t \quad (6)$$

so as to formulate a rotated DFM as a state-space model as

$$\begin{bmatrix} \mathbf{s}_t \\ m_t \end{bmatrix} = \Lambda^*(\lambda) f_t + e_t^*, \quad (7)$$

$$f_t = \mu + \Phi f_{t-1} + \eta_t \quad (8)$$

where

$$\Lambda^* = \begin{bmatrix} \beta' \\ \varphi \end{bmatrix} \Lambda$$

and

$$\Phi = I + \begin{bmatrix} \alpha_{11}^* & \alpha_{12}^* \\ \alpha_{21}^* & \alpha_{22}^* \\ \alpha_{31}^* & \alpha_{22}^* \end{bmatrix} \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & \alpha_{11}^* & \alpha_{12}^* \\ 0 & 1 + \alpha_{21}^* & \alpha_{22}^* \\ 0 & \alpha_{31}^* & 1 + \alpha_{22}^* \end{bmatrix}$$

so as to ensure that L_t is $I(1)$ and S_t and C_t are $I(0)$.

At this point, the question is: *is there an advantage in choosing representation (7) instead of (3)?* At first sight, it would seem that the two representations should be equivalent. However, any transformation of the observables in a Dynamic Factor Model has consequences on the structure of covariance matrix of the idiosyncratic shocks: the approximate DFM discards by construction the information from the off-diagonal element of the idiosyncratic shocks (e_t and e_t^* , for the two models, respectively); since it may be conjectured on theoretical grounds that these may be smaller in the rotated system, then one representation could be more efficient at picking up the signal than the other one.

Although we cannot pinpoint analytically the statistical mechanism that should lead to an improvement of the estimates, we consider it quite sensible that a representation which makes the “approximate factor model” less approximate should bring about an advantage. Moreover, the cross-correlation issue has been proven to be important for inference on the estimated factors (see Fresoli et al., 2023).

Moreover, Casoli and Lucchetti (2022) show that, in the case of cointegrated systems, a transformation that separates $I(0)$ and $I(1)$ variables can have beneficial effects on the quality of the reconstruction of the space spanned by the factors f_t . Note that, differently from Casoli and Lucchetti (2022), we do not consider m_t in differences. Instead, we assume that $m_t \sim I(1)$ and adjust the state-space model accordingly, in the spirit of Barigozzi and Luciani (2019).

In a nutshell, we extract two different sets of factors: one from the classic Nelson-Siegel model (3) and one from the rotated model (7). In the rest of the paper, we will refer to the two procedures as the “classic” and

the “rotated” method, respectively. The two equations are considered, alternatively, as the observation equation in parallel state-space models. The state transition equation for the classic model is an unrestricted VAR(1). For the rotated model, instead, we use equation (8), so as to force the level factor to be $I(1)$ and the other two factors to be $I(0)$.

4 The data

Our main analysis considers the monthly constant maturity zero-coupon yields for the major economy of the world: the US. We consider a large set of 30 maturities, ranging from 3 to 180 months (15 years). The maturities are: 3, 6, 9, 12, 15, 18, 21, 24, 27, 30, 33, 36, 42, 48, 54, 60, 66, 72, 78, 84, 90, 96, 102, 108, 120, 132, 144, 156, 168 and 180 months. We use the Treasuries data as reconstructed by Liu and Wu (2021). The dataset includes a wide set of maturities, spanning from 1 to 360 months, and samples from January 1972 to December 2024.⁸

Given that in the initial period the yields exhibit substantial volatility and that Liu and Wu (2021) evidence high pricing errors for constructing the yield curve, we decided to not to use data before January 1986. This also allows us to focus on a period that is more uniform and financially stable.

Moreover, we restrict our analysis to a set of maturities for which we are more confident about the quality of the raw data. For instance, we discard maturities shorter than 3 months because of potential noise, as well as very long maturities (i.e., more than 15 years) where the data are too sparse and the maturity distribution is characterized by relevant gaps (see Liu and Wu, 2021).

In an additional analysis, we extend our model to the European case (see Section 6 for a full description of the Euro area dataset). Figures 2 report the selected US and Euro area yields over time.

5 Results

As discussed in Section 3, we carry out estimation of the two DFM: the classic and the rotated representation, and we compare the smoothed estimates of the factors f_t from the two setups.

⁸The data are available from Jing Cynthia Wu’s website: <https://sites.google.com/view/jingcynthiawu/yield-data>.

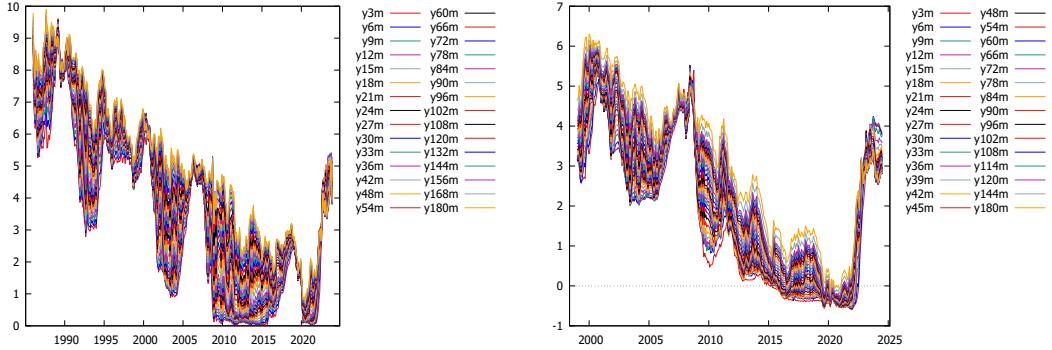


Figure 2: Monthly US (left) and European (right) yields to different maturities.

In both cases, estimation is carried out by Quasi Maximum Likelihood, assuming normality. It should be noted that, since we estimate the parameter λ rather than fixing it to a calibrated value (see subsection 5.1 below), a straightforward implementation of the EM algorithm as in Doz et al. (2012) is not easily feasible, because the elements of the loading matrix $\Lambda^*(\lambda)$ are nonlinear functions of the scalar λ . Therefore, the QML procedure employs standard numerical optimisation algorithms, such as Newton-Raphson or BFGS.

5.1 Estimation of λ

A key ingredient in estimating the state-space representation (2) is the numerical value of the scalar parameter λ , which determines the shape of the function linking the loadings to the maturities.

However, for a given value of λ , the model becomes linear and inference is much simplified. Therefore, it is very common in the literature to rely on previous findings on λ : for example, Diebold and Li (2006) perform a grid search on a range of values and set $\lambda = 0.0609$ by optimising an ad-hoc criterion. The same value that they fix is then re-used in many other papers, such as for example Inoue and Rossi (2021); Opschoor and van der Wel (2024).

In this paper, λ is instead estimated by ML along with all the other parameters in the model. By using this method, we first attempt to replicate the results of Diebold and Li (2006) and find a ML estimate for λ equal to 0.0586, with a 95% confidence interval equal to [0.0571, 0.0602]. This implies that the value used in Diebold and Li (2006) is broadly comparable in terms of orders of magnitude, but slightly outside the 95% confidence

band.

Estimating the models with our extended sample, the λ values we find are considerably smaller than the traditionally used ones: the estimates we obtain are reported in Table 2.

Method	Estimate	Std.Err.	95% CI	
classic	0.0402	0.0001	0.0401	0.0404
rotated	0.0382	0.0001	0.0379	0.0385

Table 2: Estimates of λ - US data

This result suggests that, at least if considering a sample including the late 2000s and onward, the structure of the loadings is in fact very much flatter than how often assumed by a relevant part of the literature. The value of λ set at 0.0609 is just not consistent with more recent observations, given that a relevant part of the sample includes the ZLB period.

5.2 Factor extraction

Unsurprisingly, the two methods we compare provide fairly similar results from most points of view, so we concentrate on the differences.

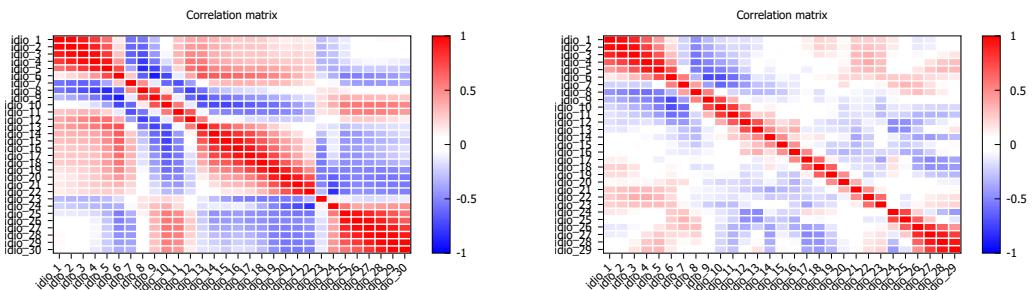


Figure 3: Correlation matrix of the idiosyncratic shocks - US data: classic method (left) vs rotated method (right)

The first difference is that the rotated model seems to perform much better at picking up the cross-correlations between the idiosyncratic disturbances of the observation equations (3) and (7). The sample correlation matrices of the idiosyncratic shocks for the two models, which are estimated as the “smoothed disturbances” from the state-space models (see Durbin and Koopman, 2012, sec. 4.5), are markedly different from one another, with the rotated model yielding a correlation matrix which is much closer to being diagonal.

In order to quantify this impression numerically, we compute the ratio

$$r_{\max} = \frac{\mu_{\max}}{n},$$

where μ_{\max} is the largest eigenvalue of the correlation matrix; clearly, this value ranges between $1/n$ and 1, with the minimum corresponding to an identity matrix and the maximum corresponding to a rank-1 matrix. This index equals 0.4014 for the classic model and 0.2264 for the rotated model. Figure 3 displays heatmaps of the two correlation matrices, from which the superiority of the rotated model is evident.

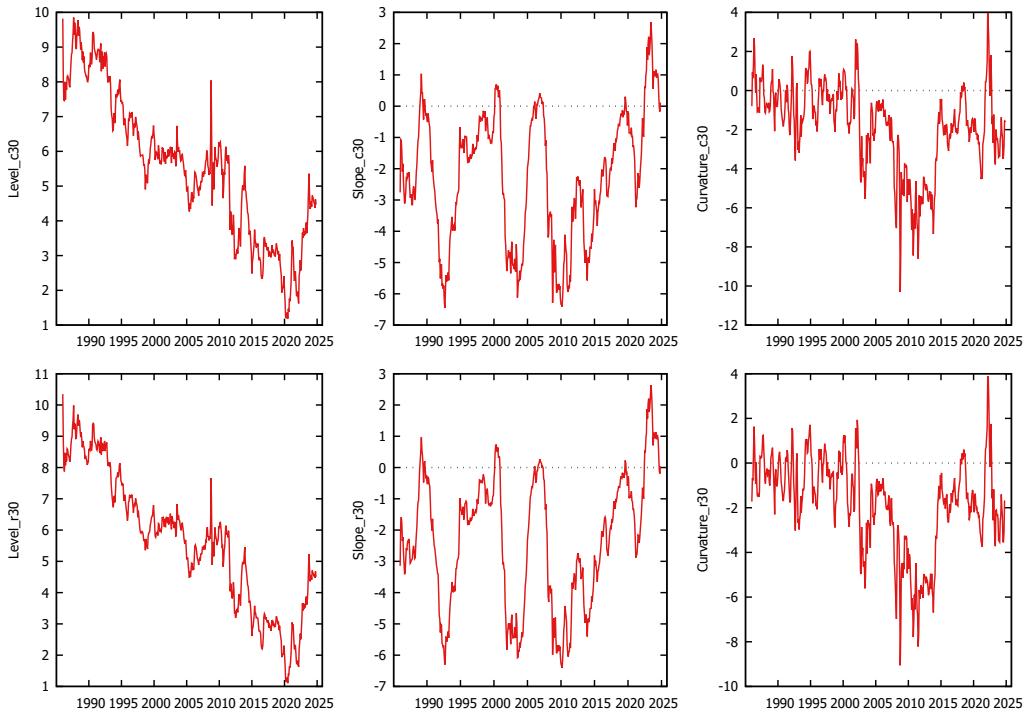


Figure 4: Estimated factors for the classic and the rotated model - US data

Another reason for preferring the rotated model comes from the observation that the correlation matrix for the classic model appears to be nearly block diagonal, with boundaries between blocks occurring at certain maturities, which is probably an artifact of the nonparametric procedure used by Liu and Wu (2021) for building the dataset and of the raw pricing data structure. Such suspicious regularity is less evident in the right-hand pane.

Figure 4 displays the estimated factors for the two methods. As anticipated, the results are visually very similar. However, the differences between the two estimates appear more evident by considering Figure 5, in

which the time path of the difference between estimated factors is shown. Interestingly, it is possible to note that the difference between the two sets of factors is larger in some periods and negligible in others. In particular, the differences become more remarkable in conjunction with economic instabilities, such for instance the period of turmoil following the US stock market crash in 1987, the early 2000s recession after the dot-com bubble, the Great Recession and the post-Covid period.

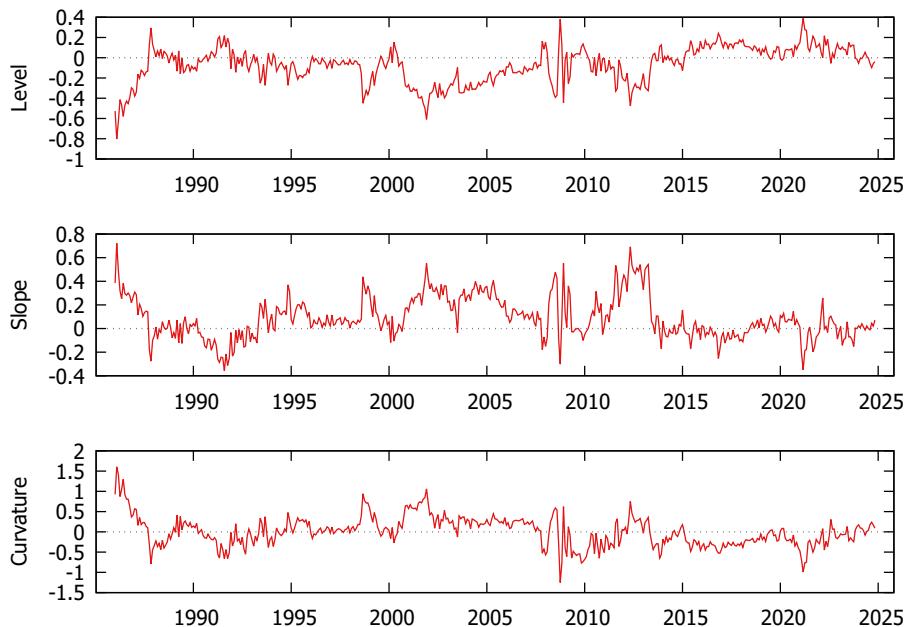


Figure 5: Differences between the factors - US data

It could be surmised that, although the estimated factors are somewhat different from one another, they in fact span exactly the same space. However, this seems not to be the case: we computed the trace R^2 index, defined as

$$TR = 1 - \frac{\text{tr} [\mathbf{E}'\mathbf{E}]}{\text{tr} [\mathbf{F}'\mathbf{F}]},$$

where \mathbf{F} is the matrix of rotated factors, \mathbf{C} is the matrix of classic factors and $\mathbf{E} = M_{\mathbf{C}}\mathbf{F}$, that is the residuals from an OLS regression of \mathbf{F} on \mathbf{C} . If the two sets spanned exactly the same space, that index would be exactly one, while in fact it equals 0.9967. Of course here the TR statistic is just used as a convenient descriptive statistic, but it is easy to see that the two sets of factors do in fact carry information that, although very similar, is partly different.

Results for a restricted set of yields, using 18 maturities instead of 30, are qualitatively similar and are available upon request.

5.3 Forecasting performance

In this subsection, we compare the performance of the factors extracted via the classic and the rotated technique for predicting a selected set of macroeconomic variables.

We begin with a preliminary remark: given the evident similarities between the two sets of factors (see Section 5.2), we do not expect dramatic differences in the forecasting power between the two methods. However, the fact that the two sets of factors are not exactly the same makes the exercise worthwhile. More specifically, we observe that known periods of economic turbulence, such as the 2007–2008 financial crisis and the subsequent downturn in economic activity, or the outbreak of the Ukraine war in 2022 seem to be reflected in the differences between the two estimates.

Therefore, we will use a comparison strategy in which we may focus on comparing the predictive power of the two factor extraction methods *at particular points in time*. In order to do so, we compare their accuracy on the basis of the log density of the forecasts (see *eg* Clements and Hendry, 1998, sec. 3.7).

The basic model we use to ascertain the forecasting power of the extracted factors is a block-triangular forecasting VAR:

$$A_{11}(L)\mathbf{z}_t = \mu_1 + A_{12}(L)\mathbf{F}_t + \varepsilon_t \quad (9)$$

$$A_{21}(L)\mathbf{F}_t = \mu_2 + \eta_t; \quad (10)$$

where \mathbf{z}_t is a vector of macroeconomic variables and \mathbf{F}_t are the estimated factors.

We concentrate on the most standard key indicators: real monthly GDP growth, the year-on-year inflation rate computed from the Consumer Price Index, and the Federal Funds effective rate.⁹

The system (9) can be thought of as a VAR model in which the absence of Granger causality from \mathbf{z}_t to \mathbf{F}_t is assumed, so that multi-step forecast is feasible given \mathbf{F}_t ; we also assume the the order of the three lag polynomials is the same and we choose it by minimising the Hannan and Quinn (1979) information criterion on the unrestricted VAR. The system is then estimated via SUR on a sample of size H , from $t - H + 1$ to $t - 1$ and

⁹All data are sourced from FRED (see <https://fred.stlouisfed.org/>). The monthly measure of GDP growth is derived in Brave et al. (2019).

the k -step-ahead forecast $\hat{\mathbf{z}}_{t+k}$ is computed, together with the associated covariance matrix $\Sigma_{t,k}$.

Given equation (9), we compute the marginal predictive likelihood by assuming normality as

$$MPL_{t,k} = -\frac{1}{2} \log [|\Sigma_{t,k}| + \mathbf{e}'_{t,k} \Sigma_{t,k}^{-1} \mathbf{e}_{t,k}]$$

where $\mathbf{e}_{t,k} \equiv \mathbf{z}_{t+k} - \hat{\mathbf{z}}_{t+k}$ is the k -step-ahead prediction error. The intuition is that a superior model should yield a forecast such that the density evaluated at the prediction error is, on average, higher.

By computing the quantity above on a rolling sample, we obtain two series of marginal predictive likelihoods, so we have MPL_t^c , in which the factors extracted with the classic method were used as \mathbf{F}_t , and MPL_t^r for our alternative (rotated) method. This indicator can be interpreted as the difference between the two models in terms of entropy:

$$PLL_{t,k} = MPL_{t,k}^r - MPL_{t,k}^c = \log \frac{\varphi(\mathbf{z}_{t+k}^r)}{f^*(\mathbf{z}_{t+k})} - \log \frac{\varphi(\mathbf{z}_{t+k}^c)}{f^*(\mathbf{z}_{t+k})}$$

where $f^*(\cdot)$ is the true unobservable density. Therefore,

$$E(PLL_{t,k}) = E \left[\log \frac{\varphi(\mathbf{z}_{t+k}^r)}{f^*(\mathbf{z}_{t+k})} \right] - E \left[\log \frac{\varphi(\mathbf{z}_{t+k}^c)}{f^*(\mathbf{z}_{t+k})} \right]$$

which can be interpreted as a (log) Kullback-Leibler divergence, where positive value indicate a better performance of the rotated model.

The approach above, therefore, makes it possible to judge the difference in forecasting power between the two approaches at any given point in time t . As argued above, however, we consider it quite important to be able to weight the differences using a metric that somehow reflects the degree of economic instability. To this aim, we adopt the approach proposed in Amisano and Giacomini (2007) and use a weighted likelihood ratio test.¹⁰ Differently from the original proposal, the weight function we use is not based on the distributional characteristics of the variable to forecast (which would be problematic anyway, given that \mathbf{z}_t is a vector), but rather on the VIX index, a commonly-used indicator of economic turmoil.

¹⁰Standard tests such as the Diebold and Mariano (1995) are not suited to compare the predictive ability in this framework for two main reasons. First, both the Diebold and Mariano (1995) and the fluctuation test of Giacomini and Rossi (2010) are designed to evaluate point forecasts, typically at the mean, rather than the entire predictive distribution. Second, these tests aggregate performance over the evaluation window, so very short episodes of instability may be smoothed out and remain undetected, as also discussed in Iacone et al. (2025).

In practice, the weight function we use is a suitable transformation of the VIX index. The results reported here use as weights a variable w_t defined as

$$\begin{aligned}\bar{V}_t &= \rho \cdot VIX_t + (1 - \rho) \cdot \bar{V}_{t-1} \\ w_t &= \frac{\bar{V}_t}{\max_{s=1 \dots T} \bar{V}_s},\end{aligned}$$

that is, a rescaling on the 0–1 interval of an exponentially weighted moving average of the VIX index (with $\rho = 1$ our weight variable is just the rescaled VIX index). We experiment with various values of ρ to reflect the fact that economic uncertainty may in fact be a smoother phenomenon than what the relatively volatile VIX index indicates.

The indicator we use for our purpose can therefore be written as

$$WLR_{t,k} = w_t \cdot PLL_{t,k} \quad (11)$$

and we check the hypothesis of equal forecasting power via the Amisano and Giacomini (2007) t -statistic, that is

$$t_k = \frac{\overline{WLR}_k}{\hat{\sigma} \sqrt{n}}$$

where $\hat{\sigma}^2$ is a heteroscedasticity and autocorrelation consistent (HAC) estimator of the asymptotic variance of $(\sqrt{n} \overline{WLR}_k)$. Under the null hypothesis of equal predictive power, this statistic has an asymptotic $N(0, 1)$ distribution.

To give the reader a pictorial example of the results, Figure 6 depicts the time paths of PLL_t , w_t and WLR_t for $k = 3$ and $\rho = 0.5$.

Table 3 displays the results of the Amisano-Giacomini test for various values of the forecasting horizon and the smoothness parameter ρ .¹¹ As it is shown, results indicate uniformly that the rotated model performs better than the classic one, with significantly better results at shorter horizons. The degree of smoothness of the weighting variable does not seem to affect the results very much.

6 The Euro Area yield curve

For the Euro area, we rely on Eikon Refinitiv data, where the Euro zero-coupon yield curve is reconstructed on a monthly basis, with maturities

¹¹We set the rolling window size to 240 months; results with a shorter window of 180 months are slightly worse, but qualitatively similar.

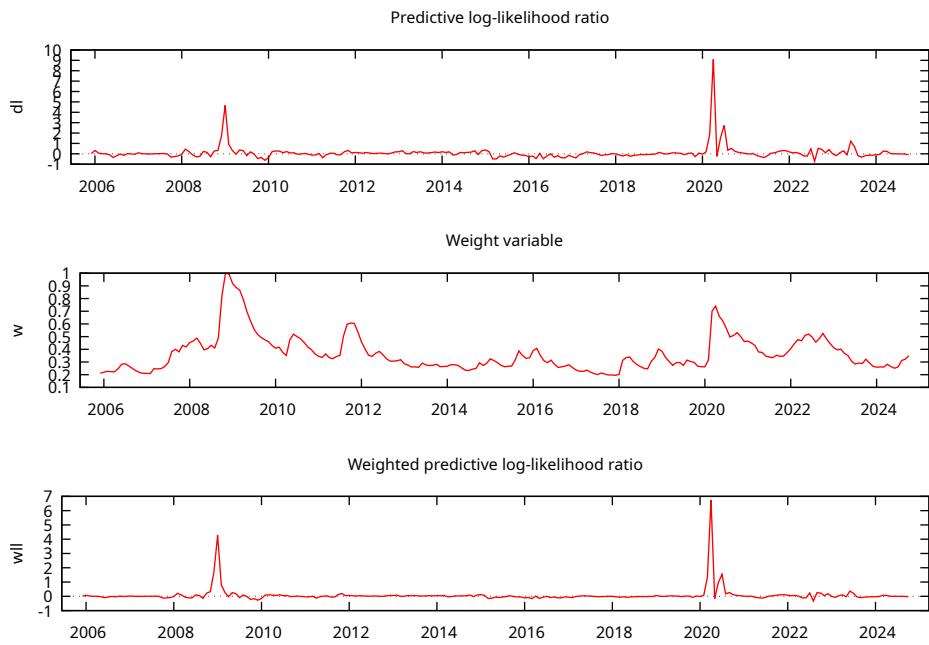


Figure 6: Amisano-Giacomini test – US data

horizon	$\rho = 0.2$	$\rho = 0.5$	$\rho = 1$
1	2.5443	2.4158	2.3924
2	2.0927	2.0377	2.0217
3	1.8438	1.8053	1.7884
6	1.7506	1.8555	1.9043
12	1.4498	1.4421	1.3326

Table 3: Amisano-Giacomini test at different k – US

ranging from 1 month to 50 years. To be consistent and provide a direct comparison in the results, we select the same maturities used for the US yield curve, spanning from 3 to 180 months. In this case, our sample starts from January 1999, that is the official start of the Euro currency.

It should be noted that, unlike the US yield curve, the European one is in fact a swap zero curve. Unfortunately, recovering a European zero-coupon curve based on risk-free bonds would be possible only by using the ECB data, which is unfeasible for two main reasons. First, yields are available only starting from 2004, making the sample's dimension not sufficient. Second, the ECB uses parametric methods to estimate the yield curve (see Svensson, 1994), which would make the data useless for our purposes: since the yield curve is reconstructed by *assuming* a factor representation, estimation of a DFM (either rotated or not) just ends up producing the same factors that were used for producing the data and, most importantly, yielding idiosyncratic errors that are almost pure numerical noise. Unfortunately, this is the case for yield curves as published by most central banks: a comprehensive, albeit not very recent list, is contained in BIS (2005).¹²

With these data, we perform an analysis similar to the one in Section 5, with a twofold aim. First, we want to make sure that our results hold for a different setting, and second, we provide evidence on a different economy to highlight similarities and disparities between the yield curve's forecasting ability in the US and the Euro area. We select, again, 30 yields, ranging from 3 to 180 months as for the US, with the difference that the selection for the Eurozone includes more short maturities and fewer long ones. Unfortunately, this choice is motivated by data availability.¹³

Most results for the US are confirmed also when analyzing European data, although the evidence for the difference between the two methods is somewhat weaker. It should be noted, however, that the dataset we have for the Euro area is considerably smaller in terms of its time span (312 monthly observations versus 468 observations for the US).

Estimated values of λ for the Euro area, shown in Table 4, resemble very much the ones for the US, with a flatter yield curve than the one implied by conventional values (in fact, even flatter than the US).

Figure 7 shows the heatmap for the idiosyncratic residuals for the model

¹²For this same reason, using other datasets such as the one provided by Gürkaynak et al. (2007) would be problematic.

¹³We carefully evaluated the possibility of selecting the same maturities for the US (i.e., picking the available yields for the Euro area and then use these also for the US), but we believe, at least when possible, it is better to pick a balanced structure including a wider set of long maturities.

Method	Estimate	Std.Err.	95% CI	
classic	0.0308	0.0002	0.0305	0.0311
rotated	0.0304	0.0002	0.0299	0.0309

Table 4: Estimates of λ – Euro data

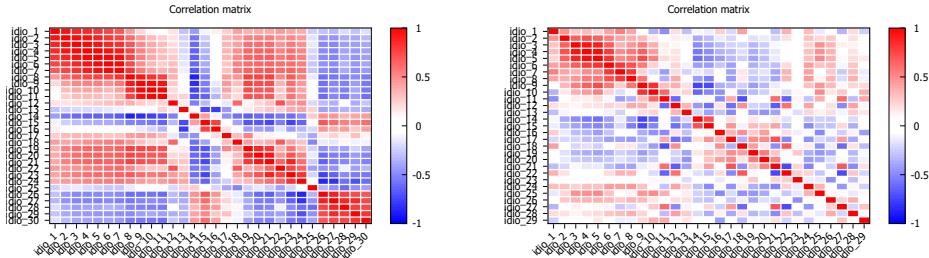


Figure 7: Correlation matrix of the idiosyncratic shocks - Euro data: classic method (left) vs rotated method (right)

with 30 maturities. The r_{max} index equals 0.498 for the classic method and 0.2989 for the rotated one. Again, there is evidence of a lesser degree of cross-correlation for the rotated model, albeit perhaps not as striking as with the US data.

A visual comparison of the estimated factors (Figures 8 and 9) leads to conclusions broadly similar to those for the US, although the peak in dissimilarity around the 2008 financial crisis is less pronounced. Other periods displaying notable divergences include the sovereign debt crisis and the recent phase of uncertainty linked to the Russia–Ukraine conflict. This is confirmed by the trace R^2 index, equal to 0.998. In the Euro area case, the similarity between the spaces spanned by the two sets of factors is much higher, although some differences are still noticeable.

For the forecasting exercise, we estimate the block-triangular VAR of Equation (9) with the inclusion of European macroeconomic variables, then focus on the Amisano-Giacomini test for assessing the predictive power of the classic and rotated models. To be consistent with the US analysis, we set a rolling window of size 240 observations.

The European macroeconomic data are taken from the ECB and OECD, selecting each time the most complete alternative in terms of observations. We consider the Hamilton-filtered Euro area monthly GDP, the year-on-year inflation rate obtained from the Harmonised Index of Consumer Prices, and the ECB's Main Refinancing Operations (MRO) rate.¹⁴

¹⁴In this case, we construct monthly GDP by temporally disaggregating the quarterly

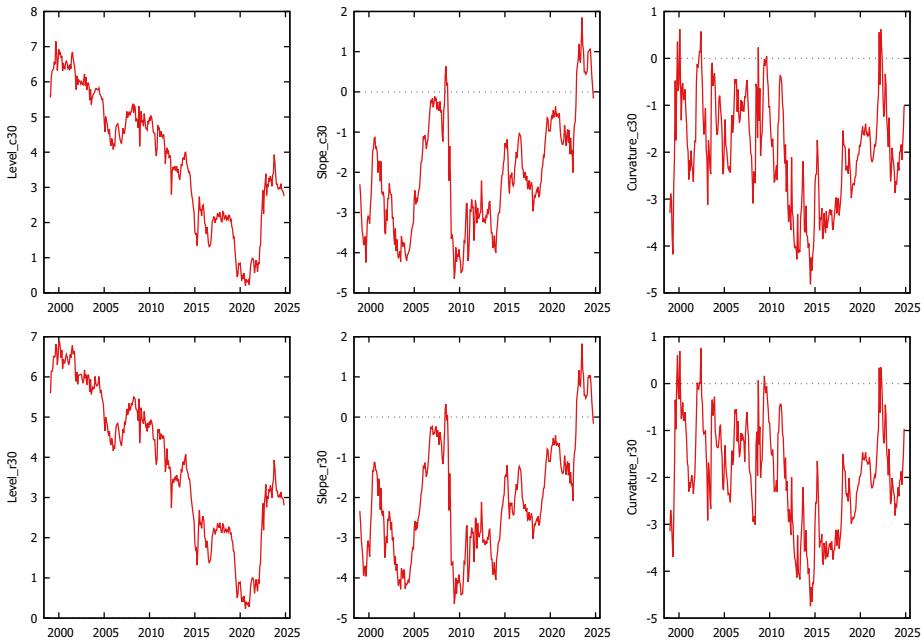


Figure 8: Estimated factors for the classic and the rotated model - Euro data

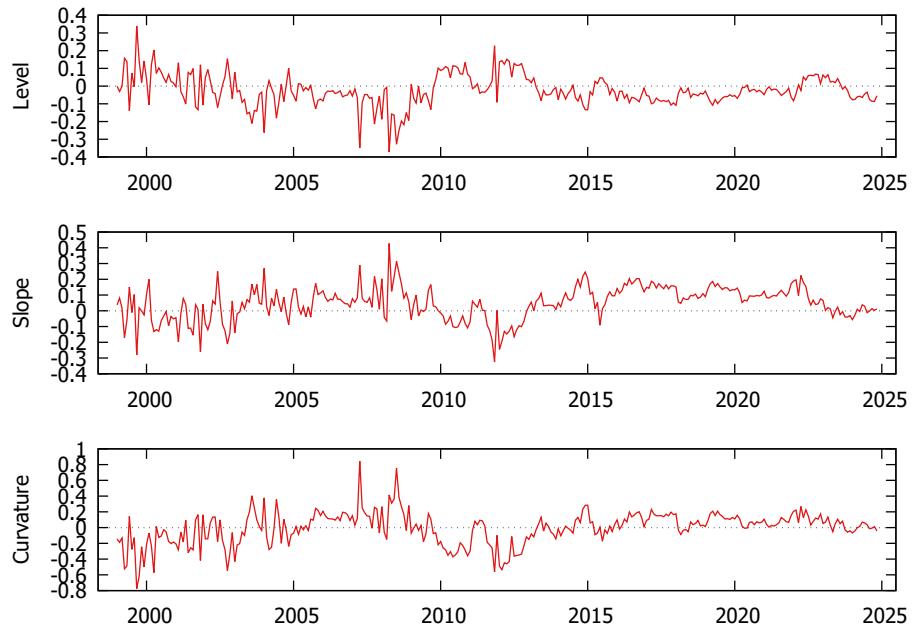


Figure 9: Differences between the factors - Euro data

series from ECB. To obtain our monthly time series, we rely on the method proposed by Fernández (1981) and use industrial production as anchor.

For the considered time span, European uncertainty and economic instability reflect a combination of financial turmoil and geopolitical shocks, unlike the US. To take this into account, we select two potential weights for the likelihood ratio test: the VSTOXX, chosen for consistency with the VIX used for the American analysis, and the Geopolitical Risk Index (GPRI) developed by Caldara and Iacoviello (2022). Since it is unclear which weight is more appropriate, we remain agnostic and report results for both, as well as for a combined weight function defined as the geometric mean of the two.

Horizon	VSTOXX		GPRI		Geometric mean	
	$\rho = 0.2$	$\rho = 1$	$\rho = 0.2$	$\rho = 1$	$\rho = 0.2$	$\rho = 1$
1	1.7359	1.7200	1.6096	1.3412	1.7112	1.6395
2	1.2292	0.9499	1.2017	0.9958	1.2435	1.0512
3	0.9397	0.5806	1.1084	1.0723	1.0460	0.8979
6	0.4898	0.0012	1.0265	0.9562	0.7884	0.5567
12	1.3097	1.0044	1.9202	2.0468	1.6659	1.6823

Table 5: Amisano-Giacomini test at different k – Euro Area

Table 5 summarises the results of the Amisano-Giacomini test for different horizons and values of ρ . In the European case, the rotated and classical forecasts produce broadly similar results, unlike in the US case. This is not surprising, as the two sets of factors for Europe are more similar. The rotated model shows statistically significant improvement at the 10% level only for 1-step-ahead forecasts with VSTOXX, and for 12-step-ahead forecasts with the GPRI index.¹⁵ For all other horizons, the differences remain not significant, although they are consistently positive, indicating a slight advantage for the rotated model.

7 Conclusions and extensions

In this article, we propose a new transformation of the Nelson-Siegel model, extended by Diebold and Li (2006), that incorporates cointegration. Our approach expresses the yield curve as a function of spreads and a common trend, and adjusts the state-space representation of the dynamic factor model accordingly. To assess the validity of our model, we compare estimates based on the original Diebold and Li (2006) representation with

¹⁵With the weight function constructed including both the series, the rotated model has a higher predicted likelihood for 1 month- and 1 year-ahead forecasts.

those obtained from our cointegration-based rotation using US yield curve data.

Our results show that, for most periods, the factors extracted by the rotated model are very similar to those from the classical specification. However, the rotation improves the model's ability to capture cross-correlations among idiosyncratic components, which translates into a more efficient extraction of the Nelson-Siegel factors during periods of financial and economic crisis. This suggests that incorporating cointegration into the model can provide additional predictive power and better reflect the information contained in the yield curve.

We then compare the forecasting ability of the two models by estimating a block-triangular VAR including standard macroeconomic variables (namely GDP growth, inflation, and the interest rate) and evaluate the differences in terms of predictive log-likelihoods. In order to do so, we rely on the Amisano and Giacomini (2007) test, which evaluates predictive accuracy at the density level. We find that the rotated model generally improves density forecasts, especially during episodes of economic stress, such as after the 2008 financial crisis and the Covid-19 recession.

The analysis on Euro area yields similar but milder improvements, thus reflecting the closer similarity between factor sets and the more limited volatility of European yields compared to the US.

Finally, we provide evidence that the parameter governing the evolution of the level, slope, and curvature factors (λ) is smaller than the value estimated in Diebold and Li (2006) and commonly used in the empirical literature. This may reflect a shift in the behavior of the term structure of interest rates following the ZLB period. This aspect, however, is not the primary focus of the present paper and will be investigated more thoroughly in further research.

Future research may also consider the application of the same techniques to economies where suitable data on the yield curve are available, and the possible gains from estimating a time-varying λ parameter.

References

Amisano G, Giacomini R. 2007. Comparing density forecasts via Weighted Likelihood Ratio tests. *Journal of Business & Economic Statistics* **25**: 177–190.

Barigozzi M, Luciani M. 2019. Quasi Maximum Likelihood estimation

of non-stationary large approximate dynamic factor models. *Papers, arXiv.org*.

BIS. 2005. Zero-coupon yield curves estimated by central banks. Technical Report 25, Bank for International settlements.

Bluwstein K, Buckmann M, Joseph A, Kapadia S, Şimşek Ö. 2023. Credit growth, the yield curve and financial crisis prediction: Evidence from a machine learning approach. *Journal of International Economics* **145**: 103773.

Bordo MD, Haubrich JG. 2024. Low interest rates and the predictive content of the yield curve. *The North American Journal of Economics and Finance* **71**: 102081.

Brand C, Cassola N. 2004. A money demand system for Euro area M3. *Applied Economics* **36**: 817–838.

Brave SA, Butters RA, Kelley D, et al. 2019. A new “big data” index of US economic activity. *Economic Perspectives, Federal Reserve Bank of Chicago* **1**.

Breitung J, Eickmeier S. 2015. Analyzing business cycle asymmetries in a multi-level factor model. *Economics Letters* **127**: 31–34.

Caldara D, Iacoviello M. 2022. Measuring geopolitical risk. *American economic review* **112**: 1194–1225.

Caldeira JF, Cordeiro WC, Ruiz E, Santos AA. 2025. Forecasting the yield curve: the role of additional and time-varying decay parameters, conditional heteroscedasticity, and macro-economic factors. *Journal of Time Series Analysis* **46**: 258–285.

Caruso A, Coroneo L. 2023. Does real-time macroeconomic information help to predict interest rates? *Journal of Money, Credit and Banking* **55**: 2027–2059.

Casoli C, Lucchetti R. 2022. Permanent-transitory decomposition of cointegrated time series via dynamic factor models, with an application to commodity prices. *The Econometrics Journal* **25**: 494–514.

Chen B, Han Y, Yu Q. 2024. Estimation and inference for CP Tensor Factor Models. *arXiv preprint arXiv:2406.17278* .

Chinn MD, Ferrara L. 2024. The predictive power of the term spread and financial variables for economic activity across countries. Technical report, National Bureau of Economic Research.

Choi I, Kim D, Kim YJ, Kwark NS. 2018. A multilevel factor model: Identification, asymptotic theory and applications. *Journal of Applied Econometrics* **33**: 355–377.

Clements M, Hendry D. 1998. *Forecasting Economic Time Series*. Forecasting economic time series. Cambridge University Press.

Coroneo L, Giannone D, Modugno M. 2016. Unspanned macroeconomic factors in the yield curve. *Journal of Business & Economic Statistics* **34**: 472–485.

Di Iorio F, Fachin S, Lucchetti R. 2016. Can you do the wrong thing and still be right? hypothesis testing in I(2) and near-I(2) cointegrated VARs. *Applied Economics* **48**: 3665–3678.

Diebold F, Li C. 2006. Forecasting the term structure of government bond yields. *Journal of Econometrics* **130**: 337–364.

Diebold F, Rudebusch G, Aruoba S. 2006. The macroeconomy and the yield curve: a dynamic latent factor approach. *Journal of Econometrics* **131**: 309–338.

Diebold FX, Mariano RS. 1995. Comparing predictive accuracy. *Journal of Business & Economic Statistics* **20**: 253–263.

Doz C, Giannone D, Reichlin L. 2012. A Quasi-Maximum Likelihood approach for large, approximate dynamic factor models. *Review of economics and statistics* **94**: 1014–1024.

Durbin J, Koopman SJ. 2012. *Time series analysis by state space methods*, volume 38. Oxford University Press.

Estrella A, Hardouvelis GA. 1991. The term structure as a predictor of real economic activity. *The journal of Finance* **46**: 555–576.

Fernández RB. 1981. A methodological note on the estimation of time series. *The Review of Economics and Statistics* **63**: 471–476.

Fresoli D, Poncela P, Ruiz E. 2023. Ignoring cross-correlated idiosyncratic components when extracting factors in dynamic factor models. *Economics Letters* **230**: 111246.

Giacomini R, Rossi B. 2010. Forecast comparisons in unstable environments. *Journal of Applied Econometrics* **25**: 595–620.

Gonzalo J, Granger C. 1995. Estimation of common long-memory components in cointegrated systems. *Journal of Business & Economic Statistics* **13**: 27–35.

Gürkaynak RS, Sack B, Wright JH. 2007. The US Treasury yield curve: 1961 to the present. *Journal of monetary Economics* **54**: 2291–2304.

Han Y, Jiao A, Ma J. 2021. The predictive power of Nelson–Siegel factor loadings for the real economy. *Journal of Empirical Finance* **64**: 95–127.

Hannan EJ, Quinn BG. 1979. The determination of the order of an autoregression. *Journal of the Royal Statistical Society. Series B (Methodological)* **41**: 190–195. ISSN 00359246.

Iacone F, Rossini L, Viselli A. 2025. Comparing predictive ability in presence of instability over a very short time. *The Econometrics Journal* : utaf018.

Inoue A, Rossi B. 2021. A new approach to measuring economic policy shocks, with an application to conventional and unconventional monetary policy. *Quantitative Economics* **12**: 1085–1138.

Kasa K. 1992. Common stochastic trends in international stock markets. *Journal of monetary Economics* **29**: 95–124.

Liu Y, Wu J. 2021. Reconstructing the yield curve. *Journal of Financial Economics* **142**: 1395–1425.

Minesso MF, Lebastard L, Le Mezo H. 2022. Text-based recession probabilities. *IMF Economic Review* **71**: 415.

Nelson CR, Siegel AF. 1987. Parsimonious modeling of yield curves. *Journal of business* : 473–489.

Opschoor D, van der Wel M. 2024. A smooth shadow-rate dynamic Nelson–Siegel Model for yields at the zero lower bound. *Journal of Business & Economic Statistics* : 1–22.

Rogoff KS, Rossi B, Schmelzing P. 2024. Long-run trends in long-maturity real rates, 1311–2022. *American Economic Review* **114**: 2271–2307.

Rossi B. 2021. Identifying and estimating the effects of unconventional monetary policy: How to do it and what have we learned? *The Econometrics Journal* **24**: C1–C32.

Svensson LE. 1994. Estimating and interpreting forward interest rates: Sweden 1992-1994.

Wahlstrøm RR, Paraschiv F, Schürle M. 2022. A comparative analysis of parsimonious yield curve models with focus on the Nelson-Siegel, Svensson and Bliss versions. *Computational Economics* : 1–38.

Wu JC, Xia FD. 2016. Measuring the macroeconomic impact of monetary policy at the zero lower bound. *Journal of Money, Credit and Banking* **48**: 253–291.

Yu R, Chen R, Xiao H, Han Y. 2024. Dynamic matrix factor models for high dimensional time series. *arXiv preprint arXiv:2407.05624* .