Erasmus University Rotterdam

Erasmus School of Economics

Bachelor Thesis Econometrics and Operations Research

Using Shifting Endpoints and a Zero Lower Bound to Forecast the Term Structure of U.S. Treasury Yields

---

Name Student:	Moneeb Akhtar
Student ID number:	543859
Supervisor:	Daan Opschoor
Second assessor:	Prof Dr. R.L. Lumsdaine

Date final version: July 3, 2022

Abstract

The yield curve could be used to determine what investors feel about the stock market and it could be used to predict recessions. Thus being able to accurately predict the yield curve is important. We investigate whether the addition of an extra factor, the addition of a lower bound and accounting for non-stationarity improves forecasting accuracy of the yields. For this, we use yield data from CRSP and Consumer Price Index data and Industrial Production data from FRED. We find that accounting for non-stationarity and adding a zero lower bound to the yield curve structure improves forecasting accuracy, however the addition of an extra factor does not.

Contents

1 Introduction

Forecasting the term structure of government bond yields is important, because yields could be used to determine what investors feel about the stock market and the economy in general. The yield curve could also be used to predict recessions, because an inverted yield curve is mostly an indicator for a recession. The behaviour of yields has changed. As an example, after the great recession and during the coronavirus pandemic, the yields were close to 0%, thus close to a lower bound. We might also need to account for potential non-stationarity of the yields which is addressed by Van Dijk et al. (2014). We thus need to apply new methods compared to the traditional methods. Therefore we are going to investigate how the yields could be forecasted in the most accurate way and our research question thus is:

Does the addition of an extra factor, accounting for non-stationarity and imposing the zero lower bound improve forecasting accuracy of the U.S. government bond yields?

We use annualized continuously compounded U.S. bond yields from June 1961 till December 2021 and we use Consumer Price Index and Industrial Production data. We apply the Dynamic Nelson-Siegel model and the Dynamic Nelson-Siegel-Svensson model with shifting endpoints, so time-varying unconditional means, and a zero lower bound to forecast the term structure of U.S. bond yields and then we compare the forecasting accuracy of these models. We find that imposing shifting endpoints improves forecasting accuracy compared to having constant endpoints. But in general, the Random Walk specification provides the most accurate forecasts. Furthermore, we find that imposing a zero lower bound generally improves forecasting accuracy, but the models with the zero lower bound do not generally give a better performance during the lower bound periods compared to periods where the yields are higher. And furthermore we find that models based on the Dynamic Nelson-Siegel framework perform better compared to models based on the Dynamic Nelson-Siegel-Svensson framework. This could help researchers to analyse and to build the most appropriate model to forecast the term structure of the U.S. government bond yields or yields of another country of region.

The remainder of this paper is as follows. In Section 2, we review previous work related to this topic. In Section 3, we discuss the data we use and in Section 4, we explain the methods. In Section 5, we present our main results and in Section 6, we conclude our paper.

2 Literature Review

To account for non-stationarity of the yields and inflation, Van Dijk et al. (2014) and Kozicki and Tinsley (2012) use shifting endpoints to forecast the yields and inflation respectively. Van Dijk et al. (2014) find that forecasting performance improves with the shifting endpoints compared to the models of Diebold and Li (2006) where they do not use shifting endpoints to account for non-stationarity of the yields. Kozicki and Tinsley (2012) find that an AR model based on shifting endpoints performs better than models with constant endpoints. That is why we introduce the subquestion:

> *Does a model with shifting endpoints deliver more accurate forecasts of the yield curve compared to models with constant endpoints?*

We expect that imposing shifting endpoints improves forecasting accuracy compared to models without those shifting endpoints, as previous research finds.

As discussed earlier, the behaviour of the yields has changed, in the sense that during the coronavirus pandemic and after the Great Recession, yields have taken values which came close to 0. To address this problem, Opschoor and van der Wel (2022) and Christensen and Rudebusch (2016) apply a zero lower bound while forecasting the U.S. bond yields. They find that the models they use with the zero lower bound incorporated in it perform better compared to the models witout the zero lower bound, especially during the periods where the interest rates are close to 0%. Therefore we introduce the subquestion:

> *Does a model with a zero lower bound provide more accurate forecasts of the yield curve compared to a model without the zero lower bound?*

We expect that models with the zero lower bound will perform better while forecasting, especially for the zero lower bound periods, as earlier research suggested. Several extensions have been considered with respect to the Nelson-Siegel model proposed by Nelson and Siegel (1987). As an example, Svensson (1995) proposed to add another factor to the Nelson-Siegel model to improve the fit of the model. Muvingi and Kwinjo (2014) use the Nelson-Siegel-Svensson model proposed by Svensson and the Nelson-Siegel model to fit the term structure of a Zimbabwean bank. They find that the Nelson-Siegel-Svensson model provides a better fit of the yield curve structure of the bank than the Nelson-Siegel model and they find that both models provide excellent forecasting abilities. Therefore we introduce the following subquestion:

> *Does a method based on the Nelson-Siegel-Svensson model provide more accurate forecasts compared to a method based on the Nelson-Siegel model?*

Based on the findings of Muvingi and Kwinjo (2014), we expect that a forecasting method based on the Dynamic Nelson-Siegel-Svensson model provides more accurate forecasts of the U.S. yields.

3 Data

We use yield data which is also used by Liu and Wu (2021). They have obtained the data from CRSP. The dataset contains the annualized continuously compounded zero coupon yields in percentage points with maturities ranging from 1 month till 360 months. The dataset covers the period between June 1961 and December 2021, thus 727 months in total. We have listed some descriptive statistics for the yields with the maturities we use in our research in months in Table 1. Based on the means of the yields for every maturity, we observe that the average yield curve is upward sloping, as the means generally increase when the maturity increases. But we also observe that the degree in which the average yields increase declines. This is thus in line with the stylized fact that the average yield curve is upward sloping and concanve. Furthermore we observe that the standard deviation generally declines when the maturity increases. Thus, that the yields are more volatile when the maturity is shorter, which is also in line with the stylized fact that the yields become less volatile as the maturity increases. For some of the maturities, our dataset does not contain information about the yields for the whole period, thus the number of observations for those maturities is lower than 727.

Table 1: Summary statistics for the US Yields between June 1961 and December 2021

Maturity	Mean	Min.	Max.	St.dev.	Obs.
3	4.50	0.01	15.95	3.26	727
6	4.66	0.03	16.13	3.31	727
9	4.77	0.04	16.11	3.33	727
12	4.84	0.06	15.96	3.33	727
18	4.97	0.11	15.94	3.32	727
24	5.06	0.12	15.72	3.29	727
36	5.24	0.12	15.57	3.22	727
48	5.40	0.17	15.48	3.16	727
60	5.51	0.23	15.20	3.08	727
84	5.72	0.38	14.95	2.98	727
96	5.94	0.44	14.94	3.15	605
108	6.01	0.49	14.94	3.11	605
120	6.07	0.53	14.94	3.04	605
180	6.39	0.73	14.91	3.00	602
240	6.09	1.05	14.78	3.00	486
360	5.28	1.29	9.48	2.01	434

For our research, we also need data containing information about inflation and industrial production growth. For that purpose, we use datasets provided by FRED, where we use the total index of the consumer price index (CPI) as a measure of inflation and the total index of industrial production (IP). For both datasets, we have monthly observations. For the CPI dataset, there are observations available from February 1947 till May 2022 at the time of writing and for the IP data we have observations between February 1919 and March 2022 at the time of writing. However, for our research, we use monthly growth rates of the two variables. We calculate the monthly growth rates according to the following equation: $growth_i = ln(x_{it}/x_{it-1})$, where $i = CPI, IP$ and x_{it} is the value of i in period t.

4 Methodology

4.1 Model specification

The first model we consider is the Dynamic Nelson Siegel-Model (DNS). The Dynamic Nelson-Siegel Model is proposed by Nelson and Siegel (1987) and Diebold and Li (2006) and is defined as

$$y_t(\tau) = \beta_{1t} + \beta_{2t}\left(\frac{1-e^{-\lambda_t\tau}}{\lambda_t\tau}\right) + \beta_{3t}\left(\frac{1-e^{-\lambda_t\tau}}{\lambda_t\tau} - e^{-\lambda_t\tau}\right) + \epsilon_t(\tau), \tag{1}$$

where $y_t(\tau)$ is the yield to maturity of a bond with maturity of τ months at time t. β_{1t}, β_{2t} and β_{3t} are the latent factors we need to estimate. The loading on β_{1t} is equal to 1 for every maturity, so β_{1t} is a long-term factor and could thus be interpreted as a level factor. The loading on β_{2t} starts at 1 when $\tau = 0$ and then declines monotonically as the maturity increases, thus β_{2t} can be interpreted as a slope factor as it is a short term factor. The loading on β_{3t} first increases when the maturity increases and then it attains its maximum value for a certain maturity, after that, the loading decreases. The factor β_{3t} is thus a curvature factor as it is more a medium term factor. Furthermore we have an error term $\epsilon_t(\tau)$, for which we assume that it has mean zero and variance σ_t^2, which does not depend on time and maturity. Diebold and Li (2006) propose a constant value of 0.0609 for λ_t, because when λ_t is set on 0.0609, it maximizes the curvature factor loading at a maturity of 30 months. We also apply this value for λ_t in our estimation and forecasting procedures.

The second model we use is based on the model proposed by Svensson (1995), where he adds another curvature factor to the Nelson-Siegel model to increase the flexibility and improve the fit of the model. We use a dynamic version of the model as Diebold and Li (2006) did with the Nelson-Siegel model. The model is defined as

$$y_t(\tau) = \beta_{1t} + \beta_{2t}\left(\frac{1-e^{-\lambda_{1t}\tau}}{\lambda_{1t}\tau}\right) + \beta_{3t}\left(\frac{1-e^{-\lambda_{1t}\tau}}{\lambda_{1t}\tau} - e^{-\lambda_{1t}\tau}\right) + \beta_{4t}\left(\frac{1-e^{-\lambda_{2t}\tau}}{\lambda_{2t}\tau} - e^{-\lambda_{2t}\tau}\right) + \upsilon_t(\tau), \quad (2)$$

where β_{4t} is an extra curvature factor and the other factors could be explained with the same reasoning as above. We name this model the Dynamic Nelson-Siegel-Svensson (DNSS) model. Christensen et al. (2009) estimate the values of λ_{1t} and λ_{2t} by using the Kalman filter on the data of the U.S. yields between January 1987 and December 2002. They find the values of 0.8379 and 0.09653 for λ_{1t} and λ_{2t} respectively. As we use U.S. bond yields as Christensen et al. (2009), we apply these values for λ_{1t} and λ_{2t} as well. Christensen et al. (2009) find these values for λ_{1t} and λ_{2t} by using maturities of $3, 6, 9, 12, 18, 24, 36, 48, 60, 84, 96, 108, 120, 180, 240$ and 360 months. Thus, for esimation and forecasting of the yields with a model based on the DNSS framework, we use the same maturities as Christensen et al. (2009). Because when doing so, the factors have their inteded meanings, so with the combination of these values for λ_{1t} and λ_{2t} and the selected maturities, we make sure that β_{1t} is a level factor, β_{2t} is a slope factor and that β_{3t} and β_{4t} are two curvature factors. With the loading on β_{3t} attaining its maximum near the 2-year maturity according to Christensen et al. (2009) and the loading on β_{4t} reaching its maximum near the 19-year maturity. To make the estimation and forecasting methods more comparable between the DNS framework and the DNSS framework, we use the same maturities as described for our estimation and forecasting for the DNS models. The difference however is, that Christensen et al. (2009) estimated their λ_{1t} and λ_{2t} by taking the maturities in years, thus we also do that when using the DNSS framework to make sure that the factors attain their intended meanings by computing the loadings based on maturities in years. For the DNS framework, we take the maturities in months to make sure that the factors have their intended meanings, as Van Dijk et al. (2014) do.

4.2 Zero Lower Bound

To address the problem of the zero lower bound (ZLB), which occured after the financial crisis of 2008 and the recent coronavirus pandemic, we impose the yield curve structure used by Opschoor and van der Wel (2022). The curve is defined as

$$\underline{y}_t(\tau) = r_{LB} + \gamma f\left(\frac{y_t(\tau) - r_{LB}}{\gamma}\right), \quad (3)$$

where r_{LB} is the lower bound for the yields. We use a lower bound of 0% as proposed by Christensen and Rudebusch (2016). And $y_t(\tau)$ is the yield curve expression as in Equation (1) and (2). $\underline{y}_t(\tau)$ is called the ZLB yield curve. Here, γ is a smoothness parameter which makes sure that the transition of the yields from the ZLB state to a state with a higher interest occurs smoothly and vice versa. For the sake of simplicity, we pre-specify γ to a value of 1. In this setup, $f(x) = x\Phi(x) + \phi(x)$, where $\Phi(\cdot)$ and $\phi(\cdot)$ are the cumulative density function and the probability density function of the standard normal distribution. For every model we consider, we forecast the yields with the ZLB restriction and without the ZLB restriction. By doing so, we are able to compare the forecasting accuracy of every model with and without the zero lower bound restriction.

4.3 Estimation

Because we fix λ_{1t} for the DNS model and λ_{1t} and λ_{2t} for the DNSS model, we can estimate the latent factors by using cross-sectional ordinary least squares at each timepoint t as Van Dijk et al. (2014). Instead of fixing those parameters, we could let them be time-varying and estimate the factors by means of nonlinear least squares, but Diebold and Li (2006) recommend the use of ordinary least squares, not only for simplicity, but also for numerical trustworthiness. As Equation (3) is a nonlinear equation, we are not able to estimate the latent factors by means of cross-sectional ordinary least squares. When we impose the zero lower bound yield curve structure on our models, we make us of nonlinear least squares (NLS) for estimation. However, $\underline{y}_t(\tau)$ is unknown when we do in-sample estimation, so for estimation, we plug the realised yields on the left hand side of the equation. And on the right hand side of the equation we plug in the DNS(S) yield curve expression instead of $y_t(\tau)$.

4.4 Forecasting

To use the Dynamic Nelson-Siegel model for forecasting, Diebold and Li (2006) define the factors as univariate AR(1) processes defined as

$$\beta_{j,t+1} = \mu_j + \phi_j(\beta_{jt} - \mu_j) + \eta_{j,t+1}. \tag{4}$$

We first forecast the factors seperately by means of Equation (4) and then plug the estimates of the factors in our models to obtain forecasts of the yields. For the DNS model, j ranges between 1 and 3, as we do have three factors. For the DNSS model, we have an extra factor, so j takes values between 1 and 4 for that model.

4.5 Shifting Endpoints

After estimating the factors for the DNS model, Van Dijk et al. (2014) find that it is not appropriate to regard the conditional mean as specified in Equation (4) to be constant. That is why they introduce a time-varying mean to account for this problem. The AR(1) process of the factors now changes to

$$\beta_{j,t+1} = \mu_{j,t+1} + \phi_j(\beta_{j,t} - \mu_{j,t}) + \eta_{j,t+1}. \tag{5}$$

This model is called the shifting endpoint model, where we now have a time-variant specification for the conditional mean which is labelled as $\mu_{j,t}$ in Equation (5). Van Dijk et al. (2014) propose various ways to estimate the means. We use the estimations by means of exponential smoothing and realised measures.

4.5.1 Exponential Smoothing

We generate the estimation of the time-varying means by means of the equation

$$\mu_{j,t+1} = \alpha\beta_{jt} + (1 - \alpha)\mu_{j,t}, \tag{6}$$

where α is a decay parameter which takes the values between 0 and 1. Here, we set α to a value of 0.1. From $t = 1$, we first iteratively compute the time-varying mean for every month. We start by setting μ_{j1} equal to β_{j1} for every factor. Then, we combine Equation (6) with Equation (5) to obtain h month ahead forecasts by iterating the process h times. We exponentially smooth all the factors in our models when we forecast based on exponential smoothing of the factors. Thus for models based on the DNS framework, we exponentially smooth the level, slope and curvature factor and for models based on the DNSS framework, we exponentially smooth the level slope and the two curvature factors.

4.5.2 Realized Measures

Diebold and Li (2006) argue that there is a correlation between inflation and the level factor and also between economic activity and the slope factor. We thus also make use of macroeconomic variables to forecast the term structure of the yields. For that purpose, we use inflation as measured by the consumer price index (CPI) and industrial production (IP). We use the monthly growth rates of the CPI and the IP as described in the Data Section. We first make use of exponential smoothing on the CPI and IP growth rates. We denote the exponentially smoothed CPI and IP as π_t^{ES} and δ_t^{ES} respectively. We follow Van Dijk et al. (2014) by setting α equal to 0.01 when we exponentially smooth our macroeconomic variables, as Van Dijk et al. (2014) argue that inflation and growth are noisy. To estimate μ_{1t}, we use the equation $\mu_{1t} = \theta_{0,1} + \theta_{1,1}\pi_t^{ES}$. For μ_{2t}, we apply the equation $\mu_{2t} = \theta_{0,2} + \theta_{1,2}\delta_t^{ES}$. The estimates of the θ variables could be obtained by applying the regressions

$$\beta_{1t} = \theta_{0,1} + \theta_{1,1}\pi_t^{ES} + \zeta_{1t} \tag{7}$$

and

$$\beta_{2t} = \theta_{0,2} + \theta_{1,2}\delta_t^{ES} + \zeta_{2t}. \tag{8}$$

To forecast the level and slope factors and thus the yields, we apply Equation (5) on each factor iteratively with the asssumption that μ_{1t} and μ_{2t} are constant at their values we obtain at the end of the sample. For every model we consider, we forecast the remaining factors via Equation (4). Thus for models based on the DNS framework, we forecast the curvature factor via Equation (4) and for models based on the DNSS framework, we forecast the two curvature factors via Equation (4).

4.6 Forecasting Methods

In Table 2 we have listed the forecasting methods we use in our research. Besides the methods with the lower bounds and the shifting endpoints incorporated, we consider a Random Walk based on the DNS and DNSS framework and the DNS and DNSS forecasting methods based on Equation (4). Table 2 only lists the methods and their acronyms without the ZLB incorporated. The acronyms of the methods with the ZLB are simply the same acronyms we use for the methods without the ZLB, but in the end we add 'LB' to them.

Table 2: Description of the methods and their acronyms

Method	Description
DNS	Dynamic Nelson-Siegel model without shifting endpoints
RW-DNS	Random walk approach with factor estimation according to DNS method
ESLSC-DNS	DNS method with exponential smoothing of all the factors
RZIG-DNS	DNS method with exponentially smoothed inflation and industrial production growth
DNSS	Dynamic Nelson-Siegel-Svensson model without shifting endpoints
RW-DNSS	Random walk approach with factor estimation according to DNSS method
ESLSC-DNSS	DNSS method with exponential smoothing of all the factors
RZIG-DNSS	DNSS method with exponentially smoothed inflation and industrial production growth

5 Results

In this Section, we present the results we have obtained. We consider two different forecasting periods. For every forecasting period, we first discuss the forecasting implementation. Then we present the results for the methods based on the DNS and DNSS frameworks seperately. And thereafter, we make a comparison between the results.

5.1 Forecasting Implementation

We use an expanding window to forecast the term structure of the interest rates. Because Christensen et al. (2009) considered the period from January 1987 to December 2002 to estimate their parameters, we use that time period as our first sample, because the λ's we use for estimation of the factors and forecasting with methods based on the DNSS methods will be close to the λ's Christensen et al. (2009) find if we start with the exact period Christensen et al. (2009) use to estimate those λ's. We forecast for all the maturities we have given in the previous Section and we forecast for three horizons ahead: 6 months, 12 months and 24 months. For all three horizons, our first forecast is made for January 2003 and the last forecast is made for December 2021. Which thus means that for every combination of forecasting method and forecasting horizon, we have in total 228 forecasts of the yields. As we make use of an expanding window, we use data back to January 1987 for every forecast we make. We compare the forecasting accuracy of the methods by computing their root mean square prediction errors (RMSPE's). While discussing the RMSPE's, we consider the maturities of 3, 12, 36, 60 and 120 months. To compare the forecasting performance of the methods with the zero lower bound to the methods without the zero lower bound when interest rates are near the zero lower bound, we make use of forecasts of the yields between January 2011 and December 2021. To estimate the parameters, we make use of data from January 1995. We consider the same maturities in our subsample as in the full sample and also the same forecasting horizons. For the subsample, we have 132 forecasts of the yields for every combination of forecasting method and forecasting horizon.

5.2 Results by Using the Full Sample

In this subsection, we provide the results we have obtained by using the full sample where we use data from 1987 to forecast the yields. First, we present the results we obtain by the methods based on the DNS framework. Then, we present the results for the methods based on the DNSS framework. After that, we compare the forecasting performance of the DNS methods with the forecasting performance of the DNSS methods.

5.2.1 Results for the Dynamic Nelson-Siegel methods

We depict the RMSPE's of the methods based on the DNS framework in Table 3. For the methods based on the Dynamic Nelson-Siegel framework, we observe that in general, the Random

Walk specification (RW-DNS) provides the smallest RMSPE's, thus provides the best forecasting accuracy when we do not take the lower bound into consideration. This is in line with what Van Dijk et al. (2014) find. Furthermore, if we compare the accuracy of the DNS method with the accuracy of the ESLSC-DNS method, we observe that in general the ESLSC-DNS method is better able to predict the term structure of the yields. For the forecasting horizon of 24 months and higher maturities, the ESLSC-DNS method provides RMSPE's which are around 40% lower compared to the RMSPE's of the DNS method. By using macroeconomic variables by means of the RZIG-DNS method, we observe that for lower forecasting horizons, its performance is worse than the perfomance of the DNS method. However, for a forecasting horizon of 24 months and especially for higher maturities when the forecasting horizon is 24 months, the RZIG method gives lower RMSPE's than the DNS method. If we compare the performance of the RZIG-DNS method to the performance of the ESLSC-DNS method, we observe that in general, the ESLSC-DNS method performs better than the RZIG-DNS method. For the horizon of 6 months, we observe that the RZIG-DNS method yields to RMSPE's which are around 20% higher compared to the RMSPE's of the ESLSC-DNS method. When imposing the lower bound on the methods, we observe that the DNS-LB method generally performs better than its counterpart without the lower bound. Sometimes, the DNS-LB method gives RMSPE's which are around 5% lower than the RMSPE's of the DNS method. When using the Random Walk, we observe that the performances of the methods with and without the lower bound are comparable to each other. The ESLSC method with the lower bound generally gives a better performance than its counterpart without the lower bound for lower maturities for every forecasting horizon. When looking at the methods with the macroeconomic variables, we observe that the lower bound method again gives more accurate forecasts in general. For the maturity of 3 months and the forecasting horizon of 6 months, the RZIG-DNS-LB method gives an RMSPE which is around 10% lower than the RMSPE the RZIG-DNS method gives. When comparing all forecasting methods in the DNS framework, we observe that the ESLSC-DNS-LB method is the method which provides the most often the lowest forecasts of all the methods we consider.

5.2.2 Results for the Dynamic Nelson-Siegel-Svensson methods

In Table 4, we show the RMSPE's for the methods based on the Dynamic Nelson-Siegel-Svensson framework with the λ's estimated by Christensen et al. (2009). For the methods based on the Dynamic Nelson-Siegel-Svensson framework, we observe that generally for these methods, the Random Walk specification, thus the RW-DNSS method yields lower RMSPE's compared to other methods based on this framework when we do not take the zero lower bound into consideration. When comparing the DNSS method to the RW-DNSS method in this case, we come to the same conclusions as before. But here we note that the RMSPE's of the DNSS method are sometimes more than twice as high as the RMSPE's of the RW-DNSS method. When forecasting the yields while using the ESLSC-DNSS method, we find that the ESLSC-DNSS method here also provides lower RMSPE's, thus better forecasts compared to the DNSS method. As in the RW-DNSS case, we also observe that the RMSPE's of the DNSS method are around twice as high compared to the ESLSC-DNSS method. Here, the RZIG-DNSS method performs better than the DNSS method for every combination of forecasting horizon and maturity. When comparing the RZIG-DNSS method to the ESLSC-DNSS method, we observe that especially for a horizon of 24 months and the maturity of 120 months, the RZIG-DNSS method provides RMSPE's which are around 50% higher compared to the RMSPE's the ESLSC-DNSS method provides. Based on the results in Table 4, we observe that in this framework, imposing the lower bound generally improves forecasting accuracy. We observe that the DNSS-LB method outperforms the DNSS method for every combination of forecasting horizon and maturity. Sometimes the DNSS-LB method improves forecasting accuracy by almost 18% compared to the DNSS method. The ESLSC-DNSS-LB method gives better forecasts than the ESLSC-DNSS method for almost every combination of maturity and forecasting horizon. However, for a maturity of 36 months and forecasting horizons of 6 and 12 months, the ESLSC-DNSS method gives better forecasts. The RZIG-DNSS-LB method also performs better than its counterpart without the lower bound, especially for lower forecasting horizons and lower maturities. Just as in the DNS case, the Random Walk specification with the lower bound gives RMSPE's which are comparable to the RMSPE's of its counterpart without the lower bound. But we observe that there is no other method which gives the lowest RMSPE's more often than the RW-DNSS-LB method, which makes this method the most accurate forecasting method when looking across all forecasting horizons and maturities in this framework, which is different to the DNS framework. There, the ESLSC-DNS-LB method was most frequently the best forecasting method.

5.2.3 Comparing the forecasting ability of each framework

If we consider the method without shifting endpoints for all frameworks, we observe that the DNS method in general yields to smaller RMSPE's compared to the DNSS method. As an example, the RMSPE of the DNSS method is 50% higher compared to the RMSPE of the DNS method when the forecasting horizon is 6 months and the maturity is 36 months. When looking at the Random Walk specification, we observe that the RMSPE's of both frameworks are comparable to each other for every combination of forecasting horizon and maturity. If there

Table 3: Root mean square prediction errors of forecasts of DNS based methods for the full sample

Forecast horizon	6 months					12 months					24 months				
Maturity	3	12	36	60	120	3	12	36	60	120	3	12	36	60	120
DNS	0.76	0.68	0.73	0.76	0.65	1.25	1.16	1.16	1.13	0.95	2.00	1.92	1.85	1.74	1.42
RW-DNS	0.68	0.63	0.62	**0.62**	0.57	1.15	1.06	0.94	0.86	0.74	1.95	1.84	**1.52**	1.28	0.95
ESLSC-DNS	0.69	0.62	**0.61**	**0.62**	**0.56**	1.20	1.06	0.93	**0.85**	**0.71**	2.15	1.94	**1.52**	**1.25**	0.90
RZIG-DNS	0.86	0.72	0.76	0.79	0.68	1.34	1.21	1.17	1.12	0.91	2.01	1.89	1.73	1.57	1.20
DNS-LB	0.70	0.65	0.70	0.74	0.65	1.18	1.11	1.12	1.11	0.94	**1.92**	1.85	1.80	1.71	1.41
RW-DNS-LB	0.67	0.63	0.62	0.63	0.57	**1.14**	1.06	0.94	0.86	0.74	1.94	1.84	**1.52**	1.29	0.96
ESLSC-DNS-LB	**0.65**	**0.60**	**0.61**	**0.62**	0.57	**1.14**	**1.04**	**0.92**	**0.85**	0.72	2.06	1.90	**1.52**	**1.25**	**0.89**
RZIG-DNS-LB	0.76	0.67	0.70	0.74	0.67	1.24	1.13	1.09	1.06	0.90	1.94	**1.83**	1.66	1.52	1.19

Note: We depict the root mean square prediction errors for forecasts of the yields. Our estimation period starts in January 1987. The first forecast of the yields is made for January 2003 and the last forecast is made for December 2021. We make use of an expanding window. The descriptions of the methods are given in Table 2. The RMPSE's in bold are the lowest RMSPE's across the methods for that combination of maturity and forecasting horizon.

are differences between the RMSPE's, they are not larger than 0.03. When comparing the methods where we exponentially smooth the means of the factors, we observe that the ESLSC-DNS method generally gives better forecasts. The ESLSC-DNSS method only gives better forecasts when the maturity of the yields is 3 months and where we take a forecasting horizon of 24 months. The differences between the forecasting accuracies between these methods with the shifting endpoints are however not as large as the differences between the methods without the shifting endpoints. We also see that the ESLSC-DNS method is more often the best performing method in its framework compared to the ESLSC-DNSS method. If we consider the approach with the macroeconomic variables, we observe that for all combinations of forecasting horizons and maturities, the RZIG-DNS method outperforms the RZIG-DNSS method. Especially for a forecasting horizon of 24 months and the maturity of 120 months, we observe that the RMSPE of the RZIG-DNSS method is 28% higher than the RMSPE of the RZIG-DNS method. When comparing the DNS-LB method to the DNSS-LB method, we come to the same conclusions as comparing these methods without the lower bounds, but the differences get smaller when incorporating the lower bound. The gains in forecasting accuracy are sometimes around 20% when imposing a lower bound on the DNSS method. But when using the DNS method, we see that the gains in accuracy are around 5%. Incorporating the lower bound on the Random Walk method does not seem to have an impact on forecasting accuracy for both frameworks, so when we compare the RW-DNSS method with the RW-DNS method, we come to the same conclusions as earlier when we compared the RW methods without the lower bound. And also for the comparison between the ESLSC-DNS-LB and ESLSC-DNSS-LB methods we can conlude the same as for the case without the lower bound, thus that the ESLSC-DNSS-LB method only gives a better RMSPE when the horizon is 24 months and the maturity is 3 months. The RZIG-DNS-LB method outperforms the RZIG-DNSS-LB method, especially for a combination of a large maturity and a large forecasting horizon, so their relative performance is also comaparable to their counterparts without the lower bound. We can conlude that in general, methods based on the DNS framework give better forecasting accuracy than methods based on the DNSS framework. The fact that generally the DNSS methods perform worse could be explained by the fact that in our forecasting framework, we forecast all the factors seperately and after forecasting those factors seperately, we plug those forecasted values for the factors in Equation (1) or (2) to obtain the forecasted yields. For the methods based on the DNS framework, we have in total 3 factors to forecast. For the DNSS methods, we need to forecast 4 factors to obtain forecasts of the yields. So for every forecasted value of the yields, we need to estimate an extra factor when forecasting with methods based on the DNSS framework, which could result in extra noise and

Table 4: Root mean square prediction errors of forecasts of DNSS based methods for the full sample

Forecast horizon	6 months					12 months					24 months				
Maturity	3	12	36	60	120	3	12	36	60	120	3	12	36	60	120
DNSS	1.13	0.95	1.10	1.22	1.21	1.82	1.70	1.82	1.89	1.86	2.75	2.70	2.74	2.72	2.62
RW-DNSS	0.68	0.64	**0.61**	**0.62**	**0.58**	1.15	**1.07**	**0.93**	**0.85**	0.75	1.94	1.86	**1.51**	**1.26**	0.98
ESLSC-DNSS	0.79	0.73	0.65	0.65	0.59	1.26	1.19	0.98	0.87	**0.74**	2.03	1.95	1.55	1.27	**0.95**
RZIG-DNSS	0.93	0.81	0.81	0.85	0.78	1.45	1.35	1.29	1.24	1.10	2.11	2.05	1.91	1.77	1.53
DNSS-LB	0.91	0.81	0.97	1.13	1.19	1.56	1.49	1.66	1.79	1.84	2.51	2.49	2.60	2.62	2.56
RW-DNSS-LB	**0.66**	**0.63**	0.62	**0.62**	**0.58**	**1.14**	**1.07**	**0.93**	**0.85**	0.75	**1.93**	**1.85**	1.52	1.27	0.98
ESLSC-DNSS-LB	0.75	0.71	0.66	0.64	**0.58**	1.23	1.16	0.99	0.87	**0.74**	2.02	1.93	1.56	1.27	**0.95**
RZIG-DNSS-LB	0.84	0.77	0.77	0.81	0.78	1.36	1.29	1.23	1.21	1.13	2.04	2.00	1.88	1.77	1.59

Note: We depict the root mean square prediction errors for forecasts of the yields. Our estimation period starts in January 1987. The first forecast of the yields is made for January 2003 and the last forecast is made for December 2021. We make use of an expanding window. The descriptions of the methods are given in Table 2. The RMPSE's in bold are the lowest RMSPE's across the methods for that combination of maturity and forecasting horizon.

less accurate forecasts of the yields and thus higher values for the RMSPE's.

5.3 Results by Using the Subsample

In this subsection, we present the results we obtain when using a subsample. We choose the subsample in such a way, that the yields for which we give forecasts are near the zero lower bound. So that means, that in our forecasting period, the period after the Great Recession and the period of the COVID-19 pandemic are primarily incorporated, as the yields are near the zero lower bound in those periods. First, we present the results of the DNS methods. Then, we do the same for the DNSS methods. Thereafter, we compare the performance of the lower bound methods in the subsample with the performance of the lower bound methods in the full sample to observe how the lower bound methods perform in the lower bound periods.

5.3.1 Results for the Dynamic Nelson-Siegel methods in the Subsample

We present the results we obtain while using the DNS based methods on our subsample in Table 5. For this sample, we again observe that when we do not consider the zero lower bound, the RW-DNS method in general provides the best forecasting results, especially for lower maturities. When comparing the RW-DNS method with the DNS method, we observe that the RW-DNS method gives lower RMSPE's for every combination of forecasting horizon and maturity, especially for the forecasting horizon of 24 months. In this sample, the ESLSC-DNS method gives RMSPE's which are generally higher when the maturities are low. But for higher maturities and especially in combination with a forecasting horizon of 24 months, the ESLSC-DNS method tends to give better results. If we compare the DNS method with the ELSC-DNS method, we see that the ESLSC-DNS method performs better for every combination of forecasting horizon and maturity, thus accounting for non-stationarity by means of shifting endpoints improves forecasting accuracy in this sample as well. The RZIG-DNS method generally performs better than the DNS method, especially for higher maturities and a higher forecasting horizon, thus the inclusion of shifting endpoints by means of macroeconomic variables also improves forecasting accuracy in this set-up. Imposing the lower bound on the DNS method generally improves forecasting accuracy, sometimes by even 30%. The performances of the DNS and DNS-LB methods are only comparable when the maturity is 120 months. In this sample, the performance of the RW-DNS-LB method also seems to be comparable to the performance of the RW-DNS method, with the differences in RMSPE's not exceeding 0.03. The ESLSC-DNS-LB method gives RMSPE's which are comparable to its counterpart without the lower bound, except for the maturity of 3 months. For every forecasting horizon, the RMSPE of the ELSC-DNS-LB method is around

10% lower compared to the RMSPE of the ESLSC-DNS method when the maturity is 3 months. For lower maturities, the RZIG-DNS-LB method tends to perform better than the RZIG-DNS method with the gain in forecasting accuracy also to be around 10%. In our subsample, we see that the RW-DNS-LB method is the method which gives the most accurate forecasts the most often. However, in our full sample, the ESLSC-DNS-LB method provides forecasts which are the best the most often. It thus seems that the performance of the methods is sensitive to the sample one chooses.

5.3.2 Results for the Dynamic Nelson-Siegel-Svensson methods in the Subsample

We show the RMSPE's obtained for the methods based on the DNSS framework in our subsample in Table 6. As is the case for other samples and frameworks, the method based on the Random Walk approach generally is the best performing method without the lower bound. However, in this set-up, the RW-DNSS method is also the method which provides the lowest RMSPE's of all methods the most frequently. Imposing shifting endpoints by means of exponential smoothing does reduce the RMSPE's in this set-up as well, with the RMSPE's of the ESLSC-DNSS method often being around twice as low as the RMSPE's of the DNS method. Generally, we do observe that the ESLSC-DNSS method gives a better forecasting accuracy than the RZIG-DNSS method, where the differences are sometimes around 10%. Only for the maturity of 3 months and the horizon of 24 months, the RZIG-DNSS method provides a smaller RMSPE. When we compare the DNS method with the DNS-LB method, we see that forecasting performance improves when incorporating the lower bound, especially for the maturity of 3 months. When the maturity is 3 months and the window is 6 months, the accuracy gain is almost 39% when adding the lower bound to the DNSS method. When adding the lower bound to the RW-DNSS method, we again observe that the lower bound generally does not change the forecasting performance. The differences between the RMSPE's stay within the margin of 0.03. The ESLSC-DNSS-LB method generally gives better forecasts than the ESLSC-DNSS method, but the differences are however never larger than 10%. For a maturity of 36 months, the ESLSC-DNSS method even outperforms the ESLSC-DNSS-LB method. Incorporating the lower bound to the RZIG-DNSS method in this subsample improves forecasting accuracy for maturities of 3 months and 12 months, but for the other maturities, the method without the lower bound gives better RMSPE's. Thus, adding the lower bound to the RZIG-DNSS method does not necessarily improve forecasting accuracy in the subsample.

5.4 Comparison of the Performance of the Lower Bound Methods Between the Samples

In this subsection, we use the Tables 3,4,5 and 6 to compare how the lower bound methods perform relative to their counterparts between the samples. In our full sample, the DNS-LB method generally improves forecasting accuracy relative to its counterpart, but the accuracy again is at most 10%. When we compare the forecasting performance between the DNS-LB method and the DNS method in our subsample, we generally see that the relative gain in forecasting accuracy is higher. As an example, for the maturity of 3 months and the horizon of 6 months, the DNS-LB method gives a RMSPE which is almost 28% lower than the RMSPE of the DNS method. For the DNSS and DNSS-LB methods, we can give the same conclusions as for the DNS case, thus the performance of the DNSS-LB method relative to its counterpart improves in the subsample. For the case where we impose shifting endpoints by means of exponential smoothing, we observe that in general the lower bound methods do better than their counterparts for both samples, however the performance of the lower bound method does not seem to improve when we consider the subsample instead of the full sample. For the ESLSC-DNSS and the ESLC-DNSS-LB methods, we generally come to the same conclusions if we compare the relative performances of the lower bound methods across samples as for the DNS case. For the methods where we use shifting endpoints by means of macroeconomic variables, we observe, as in the DNS case and the DNSS case, that the relative performance of the lower bound methods does not seem to improve when we consider the subsample relative to the full sample. We generally observe that the lower bound methods do not seem to perform better in the subsample relative to the full sample when we account for the non-stationarity of the yields. However, for the DNS and DNSS methods, where we use constant means, the lower bound methods improve forecasting accuracy in the subsample relative to the full sample, which is in line with what Opschoor and van der Wel (2022) and Christensen and Rudebusch (2016) find.

Table 5: Root mean square prediction errors of forecasts of DNS based methods for the subsample

Forecast horizon	6 months					12 months					24 months				
Maturity	3	12	36	60	120	3	12	36	60	120	3	12	36	60	120
DNS	0.68	0.50	0.62	0.74	0.76	1.08	0.94	1.08	1.19	1.17	1.63	1.57	1.76	1.84	1.76
RW-DNS	0.46	0.43	**0.46**	**0.53**	0.59	0.73	0.71	**0.71**	**0.77**	0.81	1.08	1.11	**1.09**	1.11	1.08
ESLSC-DNS	0.60	0.49	0.50	0.58	0.61	0.93	0.82	0.79	0.83	0.83	1.30	1.22	**1.09**	**1.08**	1.02
RZIG-DNS	0.61	0.56	0.57	0.66	0.66	0.94	0.91	0.91	0.95	0.88	1.26	1.27	1.25	1.25	1.12
DNS-LB	0.49	0.42	0.52	0.68	0.75	0.83	0.76	0.94	1.11	1.16	1.29	1.30	1.59	1.74	1.74
RW-DNS-LB	**0.43**	**0.41**	0.47	0.54	**0.60**	**0.72**	**0.70**	**0.71**	0.78	**0.82**	1.07	**1.09**	1.10	1.13	1.08
ESLSC-DNS-LB	0.52	0.46	0.50	0.58	0.62	0.85	0.78	0.79	0.83	0.83	1.22	1.17	**1.09**	**1.08**	**1.01**
RZIG-DNS-LB	0.52	0.51	0.52	0.61	0.66	0.83	0.82	0.84	0.91	0.89	**1.05**	**1.09**	1.16	1.21	1.12

Note: We depict the root mean square prediction errors for forecasts of the yields. Our estimation period starts in January 1995. The first forecast of the yields is made for January 2011 and the last forecast is made for December 2021. We make use of an expanding window. The descriptions of the methods are given in Table 2. The RMPSE's in bold are the lowest RMSPE's across the methods for that combination of maturity and forecasting horizon.

Table 6: Root mean square prediction errors of forecasts of DNSS based methods for the subsample

Forecast horizon	6 months					12 months					24 months				
Maturity	3	12	36	60	120	3	12	36	60	120	3	12	36	60	120
DNSS	0.93	0.71	0.91	1.09	1.14	1.45	1.30	1.50	1.66	1.71	2.11	2.03	2.19	2.27	2.30
RW-DNSS	0.47	0.44	**0.46**	**0.54**	**0.60**	0.74	0.72	**0.71**	**0.77**	**0.82**	1.07	1.12	1.08	1.11	1.10
ESLSC-DNSS	0.51	0.49	0.49	0.59	0.63	0.77	0.78	0.74	0.81	0.84	1.08	1.14	**1.07**	**1.07**	1.06
RZIG-DNSS	0.49	0.44	0.56	0.75	0.83	0.73	0.71	0.85	1.03	1.14	1.03	1.08	1.23	1.34	1.45
DNSS-LB	0.53	0.46	0.64	0.88	1.05	0.95	0.92	1.19	1.44	1.60	1.69	1.70	1.95	2.11	2.20
RW-DNSS-LB	0.44	0.42	0.47	0.55	**0.60**	0.72	0.70	0.72	0.78	**0.82**	1.06	1.10	1.10	1.13	1.09
ESLSC-DNSS-LB	0.46	0.46	0.50	0.57	0.62	0.74	0.75	0.76	0.80	0.83	1.08	1.12	1.09	1.08	**1.05**
RZIG-DNSS-LB	**0.42**	**0.41**	0.52	0.74	0.91	**0.65**	**0.66**	0.83	1.09	1.32	**0.93**	**0.98**	1.24	1.48	1.73

Note: We depict the root mean square prediction errors for forecasts of the yields. Our estimation period starts in January 1995. The first forecast of the yields is made for January 2011 and the last forecast is made for December 2021. We make use of an expanding window. The descriptions of the methods are given in Table 2. The RMPSE's in bold are the lowest RMSPE's across the methods for that combination of maturity and forecasting horizon.

6 Conclusions

For our reseach, we try to answer the question: *"Does the addition of an extra factor, accounting for non-stationarity and imposing the zero lower bound improve forecasting accuracy of the U.S. government bond yields?"*. We observe that in general, the methods where we account for non-stationarity of the yields give a better forecasting accuracy of the yields compared to methods where we do not account for the non-stationarity of the yields, so the DNS and DNSS methods. However, the Random Walk specifications tend to give the best forecasting results. These findings also answer our first subquestion whether the addition of shifting endpoints improves forecasting accuracy.

Furthermore, we use the zero lower bound yield curve structure proposed by Opschoor and van der Wel (2022) to consider whether the addition of the lower bound improves forecasting accuracy. We conclude that imposing a lower bound improves the overall forecasting accuracy. But for methods where we already account for the non-stationarity of the yields, we observe that after the addition of the lower bound, the forecasting performance does not improve during the lower bound periods compared to periods where the yields are generally higher. This also answers our second subquestion whether the addition of the lower bound structure improves forecasting accuracy or not.

To answer our third subquestion whether a method based on the DNSS framework gives a better forecasting performance than a method based on the DNS framework, we consider methods based on the DNSS framework relative to the DNS methods to see whether the addition of an extra factor improves forecasting accuracy. Generally we find that methods based on the DNS framework yield to better forecasts of the U.S. bond yields compared to methods based on the DNSS framework, so the addition of an extra factor does not improve forecasting accuracy, which is not in line with what we expected. But this result could be explained by the fact that more factors need to be forecasted seperately when using the DNSS method, which could result in more noise and less accurate forecasts. This also answers our research question. The addition of a lower bound and accounting for non-stationarity improves forecasting accuracy, but the addition of another factor does not.

For further research we could use different estimation techniques of the yield curve structure for forecasting purposes. We could perhaps estimate the yields by means of the Kalman filter as done by Opschoor and van der Wel (2022) and Christensen et al. (2009). We could also perhaps compare the forecasting performance of the methods we consider by using the yields of other countries, such as European countries.

References

Christensen, J. H., Diebold, F. X., & Rudebusch, G. D. (2009). An arbitrage-free generalized nelson–siegel term structure model.

Christensen, J. H., & Rudebusch, G. D. (2016). Modeling yields at the zero lower bound: Are shadow rates the solution? *Dynamic factor models*. Emerald Group Publishing Limited.

Diebold, F. X., & Li, C. (2006). Forecasting the term structure of government bond yields. *Journal of econometrics*, *130*(2), 337–364.

Kozicki, S., & Tinsley, P. A. (2012). Effective use of survey information in estimating the evolution of expected inflation. *Journal of Money, Credit and Banking*, *44*(1), 145–169.

Liu, Y., & Wu, J. C. (2021). Reconstructing the yield curve. *Journal of Financial Economics*, *142*(3), 1395–1425.

Muvingi, J., & Kwinjo, T. (2014). Estimation of term structures using nelson-siegel and nelson-siegel-svensson: A case of a zimbabwean bank. *Journal of Applied Finance and Banking*, *4*(6), 155.

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

Opschoor, D., & van der Wel, M. (2022). A smooth shadow-rate dynamic nelson-siegel model for yields at the zero lower bound.

Svensson, L. E. (1995). Estimating forward interest rates with the extended nelson & siegel method. *Sveriges Riksbank Quarterly Review*, *3*(1), 13–26.

Van Dijk, D., Koopman, S. J., Van der Wel, M., & Wright, J. H. (2014). Forecasting interest rates with shifting endpoints. *Journal of Applied Econometrics*, *29*(5), 693–712.

A Appendix

A.1 Code Description

To obtain our results, we use *R* version *4.2.0* and the programme *RStudio* with version *2022.02.2+485*. In total, we use 4 files for our analysis. In the file "DNS_DNSS", we forecast the yields using our full sample starting from January 1987 by making use of the DNS(S),RW-DNS(S),ESLSC-DNS(S) and RZIG-DNS(S) methods. In the file "DNSZLB_DNSSZLB", we also make use of the full sample starting from January 1987 and we use the methods described above with the lower bound to forecast the yields. Here, we make use of nonlinear least squares to estimate the factors. For estimation, we make use of starting values in our nonlinear estimation which are equal to 1 for all the factors. We forecast the yields for the period between January 2003 and December 2021.

To forecast the yields in our subsample without the lower bound incorporated, we make use of the file "DNS_DNSS_GR_COVID". We again consider the methods we have mentioned above, but this time our estimation starts from January 1995. and the yields are forecasted for the period between January 2011 and December 2021. In the file "DNSZLB_DNSSZLB_GR_COVID", we consider the same methods with the zero lower bound. Our estimation and forecasting samples are same as for the "DNS_DNSS_GR_COVID" file. Because we make use of the zero lower bound, we estimate the factors by means nonlinear squares. When doing so, we again use starting values equal to 1 for all the factors.