ARTICLE IN PRESS
REVECO-00803; No of Pages 15
International Review of Economics and Finance xxx (2013) xxx–xxx
Contents lists available at SciVerse ScienceDirect
International Review of Economics and Finance
journal homepage: www.elsevier.com/locate/iref

Estimating the spot rate curve using the Nelson–Siegel model A ridge regression approach

Jan Annaert [a,b], Anouk G.P. Claes [c], Marc J.K. De Ceuster [a,b,*], Hairui Zhang [a,b]

[a] *Universiteit Antwerpen, Prinsstraat 13, 2000 Antwerp, Belgium*
[b] *Antwerp Management School, Sint-Jacobsmarkt 13, 2000 Antwerp, Belgium*
[c] *Université Saint-Louis Brussels & Louvain School of Management Group, Boulevard du Jardin Botanique 43, 1000 Brussels, Belgium*

ARTICLE INFO

Article history:
Received 9 January 2012
Received in revised form 4 January 2013
Accepted 9 January 2013
Available online xxxx

JEL classification:
E43
C51

Keywords:
Smoothed bootstrap
Ridge regression
Nelson–Siegel
Spot rates

ABSTRACT

The Nelson–Siegel model is widely used in practice for fitting the term structure of interest rates. Due to the ease in linearizing the model, a grid search or an OLS approach using a fixed shape parameter are popular estimation procedures. The estimated grid search parameters, however, have been reported (1) to behave erratically over time, and (2) to have relatively large variances. On the other hand, parameter estimates based on a fixed shape parameter, while avoiding multicollinearity, turn out to be too smooth. We show that the Nelson–Siegel model can become heavily collinear depending on the estimated/fixed shape parameter. A simple procedure based on ridge regression can remedy the reported problems significantly.

1. Introduction

Good estimates of the term structure of interest rates are of the utmost importance to investors and policy makers. One of the term structure estimation methods, initiated by Bliss and Fama (1987), is the smoothed bootstrap which bootstraps discrete spot rates from market data and then fits a smooth and continuous curve to the data. Although various curve fitting spline methods have been introduced (quadratic splines by McCulloch, 1971; cubic splines by McCulloch, 1975; exponential splines by Vasicek & Fong, 1982; B-splines by Shea, 1984), these methods have been criticized on the one hand for having undesirable economic properties and on the other hand for being 'black box' models (Seber & Wild, 2003). Nelson and Siegel (1987) and Svensson (1994, 1996) therefore suggested parametric curves that are flexible enough to describe a whole family of observed term structure shapes. These models are parsimonious, consistent with a factor interpretation of the term structure (Litterman & Scheinkman, 1991) and have both been widely used in academia and in practice. In addition to the level, slope and curvature components present in the Nelson–Siegel (NS) model, the Svensson model contains a second hump/trough term which allows for an even broader and more complicated range of term structure shapes. In this paper, we restrict ourselves to the NS model. The Svensson model shares – by definition – all the reported problems of the NS approach. Since the source of the problems, i.e., collinearity, is the same for both models, the reported estimation problems of the Svensson model may be reduced analogously.

The NS model is extensively used by central banks and monetary policy makers (Bank of International Settlements, 2005; European Central Bank, 2008). Fixed-income portfolio managers use the model to immunize their portfolios (Barrett, Gosnell, &

* Corresponding author at: Universiteit Antwerpen, Prinsstraat 13, 2000 Antwerp, Belgium.
E-mail address: marc.deceuster@ua.ac.be (M.J.K. De Ceuster).

1059-0560/$ – see front matter © 2013 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.iref.2013.01.005

Heuson, 1995; Hodges & Parekh, 2006) and recently, the NS model also regained popularity in academic research. Dullmann and Uhrig-Homburg (2000) use the NS model to describe the spot rate curves of Deutsche Mark-denominated bonds to calculate the risk structure of interest rates. Fabozzi, Martellini, and Priaulet (2005) and Diebold and Li (2006) benchmarked the NS forecasts against other models in the term structure forecasts, and found it performs well, especially for longer forecast horizons. Martellini and Meyfredi (2007) use the NS approach to calibrate the spot rate curves and estimate the value-at-risk for fixed-income portfolios. Yan, Shi, and Wu (2008) use the NS model to bootstrap riskless spot rate curve as the input for calculating the credit risk spread for the U.S. market. Finally, the NS model estimates are also used as an input for affine term structure models. Coroneo, Nyholm, and Vivada-Koleva (2011) test to which degree the NS model approximates an arbitrage-free model. They first estimate the NS model and then use the estimates to construct interest rate term structures as an input for arbitrage-free affine term structure models. They find that the parameters obtained from the NS model are not statistically different from those obtained from the 'pure' no-arbitrage affine-term structure models.

Notwithstanding its economic appeal, the NS model is highly nonlinear which causes many users to report estimation difficulties. Nelson and Siegel (1987) transformed the nonlinear estimation problem into a simple linear problem, by fixing the shape parameter that causes the nonlinearity. In order to obtain parameter estimates, they computed the ordinary least squares (OLS) estimates of a series of models conditional upon a grid of the fixed shape parameter. The estimates that, conditional upon a fixed shape parameter, maximized the R^2 were chosen. We refer to their procedure as a *grid search*. Others have suggested estimating the NS parameters simultaneously using *nonlinear optimization techniques*. Cairns and Pritchard (2001), and Vermani (2012), however, show that the estimates of the NS model are very sensitive to the starting values used in the optimization. Moreover, the time series of the estimated coefficients have been documented to be very unstable (Barrett et al., 1995; de Pooter, 2007; Diebold & Li, 2006; Fabozzi et al., 2005; Gurkaynak, Sack, & Wright, 2006) and even to generate negative long term rates, thereby clearly violating any economic intuition. Finally, the standard errors on the estimated coefficients, though seldom reported, are large.

Although these estimation problems have been recognized before, it has never led towards satisfactory solutions. Instead, it became a common practice to fix the shape parameter over the whole time series of the term structures.[1] Hurn, Lindsay, and Pavlov (2005), however, point out that the NS model is very sensitive to the choice of this shape parameter. de Pooter (2007) confirms this finding and shows that with different fixed shape parameters, the remaining parameter estimates can take extreme values. We show that fixing the shape parameter can also result in extremely smooth time series of the parameter estimates, making it a non-trivial issue. To alleviate the observed problems substantially and to estimate the shape parameter freely, we use ridge regression.

The remainder of this paper is organized as follows. In Section 2, we introduce the NS model. Section 3 presents the estimation procedures used in the literature, illustrates the multicollinearity issue which is conditional on the estimated (or fixed) shape parameter and proposes an adjusted procedure based on the ridge regression. In the subsequent section (Section 4) we present our data and their descriptive statistics. Since the ridge regression introduces a bias in order to avoid multicollinearity, we mainly evaluate the merits of the models based on their ability to extrapolate the short and long end of the term structure. The estimation results and the robustness of our ridge regression are discussed in Section 5. Finally, we conclude.

2. A first look at the Nelson–Siegel model

The Nelson and Siegel (1987) spot rate function $r(\tau)$ at time to maturity τ is specified as

$$r(\tau)=\begin{bmatrix}\beta_0\\ \beta_1\\ \beta_2\end{bmatrix}'\begin{bmatrix}1\\ \lambda\left(1-e^{-\tau/\lambda}\right)/\tau\\ \lambda\left(1-e^{-\tau/\lambda}\right)/\tau-e^{-\tau/\lambda}\end{bmatrix}=\begin{bmatrix}\beta_0\\ \beta_1\\ \beta_2\end{bmatrix}'\begin{bmatrix}r_0\\ r_1\\ r_2\end{bmatrix}. \tag{1}$$

In Eq. (1) r_0, r_1 and r_2 represent the level, slope and curvature components of the spot rate curve. The role of the components becomes clear when we look at their limiting behavior with respect to the time to maturity. When the time to maturity grows to infinity, the slope and curvature component vanish and the long-term spot rate converges to a constant level of interest rate, β_0. When the time to maturity approaches zero, only the curvature component vanishes and the spot rate converges to $(\beta_0+\beta_1)$. The spread, $-\beta_1$, measures the slope of the term structure, whereby a negative (positive) β_1 represents an upward (downward) slope. The degree of the curvature is controlled by β_2, the rate at which the slope and curvature component decay to zero. Finally, the location of the maximum/minimum value of the curvature component is determined by λ. Note that λ determines both the shape of the curvature component and the hump/trough of the term structure. By maximizing the curvature component in the spot rate function with respect to λ, we are able to determine the location of the hump/trough of the term structure. The curvature component in the spot rate curve reaches its maximum when $\tau>\lambda$, which is determined by simply maximizing r_2 in Eq. (1) with λ fixed. Alternatively, we can force the location of the hump/trough of the term structure to be at a given time to maturity, by fixing the shape parameter to a specific value. This also linearizes the model and hence facilitates estimation (as in Diebold & Li, 2006; Fabozzi et al., 2005).

[1] Barrett et al. (1995) and Fabozzi et al. (2005) fix this shape parameter to 3 for annualized returns. Diebold and Li (2006) choose an annualized fixed shape parameter of 1.37 to ensure the stability of parameter estimation.

3. Estimation procedures

Taking the drawbacks of the nonlinear regression techniques into account (see e.g., Cairns & Pritchard, 2001; Ferguson & Raymar, 1998; Vermani, 2012), most researchers fixed the shape parameter and estimated a linearized version of the NS model. The parameters of the NS model have typically been estimated by minimizing the sum of squared errors (SSE) by using (1) OLS over a grid of pre-specified λ's (Nelson & Siegel, 1987), and (2) a linear regression, conditional on a chosen fixed shape parameter λ (de Pooter, 2007; Diebold & Li, 2006; Fabozzi et al., 2005). We refer to these methods as the traditional measures. The estimated parameters using the traditional methods, however, are reported to behave erratically in time, and to have relatively large variances. Alternatively, parameter estimates based on a fixed shape parameter, while avoiding multicollinearity, will prove to be too smooth. We introduce a ridge regression approach to remedy the reported problems and we show how to judiciously fix the shape parameter in order to estimate the linearized NS model.

3.1. The nature of the multicollinearity problem

Researchers are aware of potential multicollinearity issues when estimating the NS model. Diebold and Li (2006), for example, indicate that the high correlation between the slope and the curvature component of the NS model makes it difficult to obtain good estimates of the parameters. What seems to have gone unnoticed, however, is the fact that the correlation between the two regressors of the model depends on (the times to maturity of) the financial instruments chosen in the bootstrap. In order to illustrate this point, we consider four different sets of times to maturity: *Vector 1*: 3 and 6 months, 1, 2, 3, 4, 5, 7, 10, 15, 20 and 30 years; *Vector 2*: 3, 6, 9, 12, 15, 18, 21, 24, 30 months, 3–10 years; *Vector 3*: 1 week, 1–12 month, and 2–10 years; and *Vector 4*: 1 week, 6 months, and 1–10 years. The first set of maturities was used by Fabozzi et al. (2005). Diebold and Li (2006) opted for the second maturity vector. Since researchers are inclined to use all the data that they can find, we study two extra vectors that also include additional shorter time to maturities.

We calculate the correlations between the regressors for the λ values chosen by Diebold and Li (2006) and Fabozzi et al. (2005), by using the four vectors of time to maturity that we consider. When λ is fixed to 1.37 (Diebold & Li, 2006) the correlations are 0.256; −0.051; −0.549 and −0.352 for the 4 respective maturity vectors. Fixing λ to 3 as in Fabozzi et al. (2005), generates more negative correlations: they become −0.324, −0.871, −0.931 and −0.872 for vectors 1 to 4. This shows that the correlation between the slope and the curvature components of the NS model heavily depends on the choice of the shape parameter. The second maturity vector e.g., implies correlations varying between −5% and −87% depending on the chosen shape parameter. The correlation also severely depends on the choice of the time to maturity vector. Using $\lambda = 1.37$, the correlation varies from −0.549 to 0.256, for the maturity vectors chosen. Setting $\lambda = 3$ produces correlations from −0.324 to −0.931. The vector containing the series of short maturities (the third vector) turns out to be the most sensitive to the collinearity issue.

Fig. 1 gives a more complete picture by plotting the correlation between the two regressors over a range of λ values using the four time to maturity vectors studied. We notice that the choice of the maturity vector influences the steepness of the correlation curve. It appears that Fabozzi et al. (2005) and Diebold and Li (2006) chose their λ values judiciously conditional on their maturity vector, although they both motivate their choice differently. It is clear that for empirical work it is of the utmost importance for all of the estimation methods to take this potential multicollinearity issue into account.

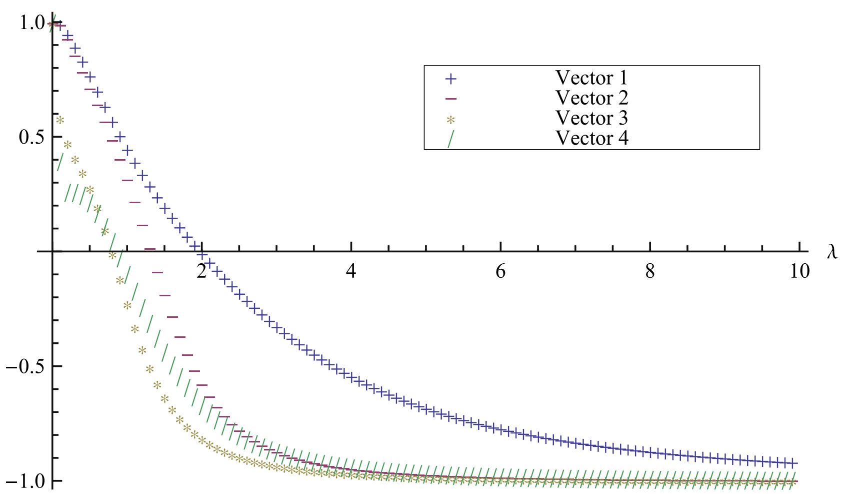

Fig. 1. Correlations between the slope and curvature components.

Please cite this article as: Annaert, J., et al., Estimating the spot rate curve using the Nelson–Siegel model, *International Review of Economics and Finance* (2013), http://dx.doi.org/10.1016/j.iref.2013.01.005

3.2. Traditional estimation methods

3.2.1. Grid search based OLS

To avoid nonlinear estimation procedures, Nelson and Siegel (1987) linearize their model by fixing λ and estimate Eq. (1) with OLS. This procedure was repeated for a whole grid of λ values ranging from 0.027 to 1. The estimates with the highest R^2 were then chosen as the optimal parameter set. In practice, it is well known that a grid search based OLS leads to parameter instability in the time series of estimates. This has been pointed out by many researchers including Barrett et al. (1995), Cairns and Pritchard (2001), Fabozzi et al. (2005), Diebold and Li (2006), Diebold, Li, and Yue (2008), Gurkaynak et al. (2006) and de Pooter (2007). Far less observed is that extreme multicollinearity among the two regressors is the source of this instability. Moreover, high multicollinearity can also inflate the variance of the estimators.

3.2.2. OLS with fixed shape parameters

Some researchers fix the shape parameter which they typically motivate by prior knowledge about the curvature of the spot rates.

- Diebold and Li (2006) set λ to 16.4 with monthly compounded returns, or approximately 1.37 with annualized data. Their choice implies that the curvature component in the spot rate function will have its maximal value at the time to maturity of 2.5 years. They motivated their choice by stating that most of the humps/troughs are between the second and the third years. As we have shown, this choice also turned out to avoid multicollinearity problems for their maturity vector.
- Following Barrett et al. (1995), Fabozzi et al. (2005) fixed λ to 3 with annualized data, implying the hump to be located at a maturity of 5.38. Barrett et al. (1995) performed a grid search by fixing λ for the whole dataset and obtained a global optimal shape parameter in one go. In a footnote, Fabozzi et al. (2005) mention that when the shape parameter is fixed at 3, the correlation between the two regressors will not cause severe problems for their data, but do not offer a universal procedure to tackle the multicollinearity issue.

From the previous section, we know that the degree of multicollinearity depends on both the choice of the fixed shape parameter and the choice of the vector of times to maturity. The shape parameter that minimizes the squared errors may also vary over time.

3.2.3. Grid search with conditional ridge regression

Whereas linear regressions do not require starting values for the estimators and always give globally optimal estimators, they do suffer from instability in parameter estimation. In this paper, we follow Nelson and Siegel's (1987) approach by combining the grid search with the OLS regression to 'free' the shape parameter. Conditional on the λ that results in the highest R^2, the parameters are re-estimated by using ridge regression whenever the degree of multicollinearity among the regressors is 'too high'. We therefore need to test the degree of multicollinearity of the two regressors. The measure we use is discussed below. Subsequently, we discuss the nature of the ridge regression and present the implementation of the ridge regression for the NS term structure estimation.

3.2.3.1. Measuring the degree of multicollinearity. In order to address the multicollinearity issue, we need to verify the degree of collinearity. Here we follow e.g., Haile and Pozo (2008) and use the condition number as the collinearity measure.[2] Assume there is a standardized linear system $\mathbf{y} = \boldsymbol{\beta}\boldsymbol{X} + \boldsymbol{\varepsilon}$. Denote κ (kappa) as the condition number and ν the eigenvalues of $\mathbf{X}'\mathbf{X}$. The condition number of $\mathbf{X}$ is defined as

$$\kappa(\mathbf{X}) = \frac{\nu_{\max}}{\nu_{\min}} \geq 1. \tag{2}$$

If $\mathbf{X}$ is well-conditioned (i.e., the regressor columns are uncorrelated), then the condition number is one, which implies that the variance is explained equally by all the regressors. If correlation exists, then the condition number is no longer equal to 1. The difference between the maximum and minimum eigenvalues will grow as the collinearity effect increases. As suggested by Belsley (1991), we use a condition number of 10 as a measure of the degree of multicollinearity.[3]

3.2.3.2. Remedy of high collinearity. Once collinearity is detected, we need to remedy the problem. To overcome the OLS parameter instability due to multicollinearity, we implement ridge regression. This estimation procedure can substantially reduce the sampling variance of the estimator, by adding a small bias to the estimator. Kutner, Nachtsheim, Neter, and Li (2004) show that biased estimators with a small variance are preferable to unbiased estimators with a large variance, because the small variance estimators are less sensitive to the measurement errors. We therefore use the ridge regression and compute our estimates as follows:

$$\widehat{\beta}^* = \left[\mathbf{X}'\mathbf{X} + k\mathbf{I}\right]^{-1}\mathbf{X}'\mathbf{y}, \tag{3}$$

where k is called the ridge constant, which is a small positive constant.

[2] Alternatively we could have used the variance inflation factor (VIF) as in e.g., Lee, Xie, and Yau (2011). But Belsley (1991) pointed out that a high VIF is a sufficient but not a necessary condition for a collinearity problem. Therefore we prefer to base our inference on the condition number.

[3] As there are only two regressors in the NS model, we can plot a one-to-one relationship between the condition number and the correlation between the two regressors. In our dataset, a condition number above 10 is equivalent to a correlation with an absolute value above 0.8.

As mentioned by Kutner et al. (2004), ridge regression estimates are obtained by using the method of penalized least squares, which minimizes the combination of the sum of squared errors with a roughness penalty function:

$$Q = SSE + k\left[\sum_{j=1}^{p}\left(\widehat{\beta}_j^*\right)^2\right], \tag{4}$$

where $\widehat{\beta}^* = \left[\widehat{\beta}_1^*, \widehat{\beta}_2^* \cdots \widehat{\beta}_p^*\right]$, and SSE is the sum of squared errors. We can see that the higher the coefficient is, the higher the penalty on it will be. As a result, the ridge regression based estimates have a smaller magnitude than those based on the OLS. As the ridge constant increases, the bias grows and the estimator variance decreases, along with the condition number. Clearly, when $k = 0$ the ridge regression is a simple OLS regression.

3.2.3.3. Implementation. As pointed out by Kutner et al. (2004), collinearity increases the variance of the estimators and makes the estimated parameters unstable. However, even under a high collinearity, the OLS regression still generates the unbiased estimates. As a result, we implement a combination of the grid search and the ridge regression by using the following steps:

1. Perform a grid search based on the OLS regression to obtain the estimate of λ which generates the lowest mean squared error.
2. Calculate the condition number for the 'optimal' λ.
3. Re-estimate the coefficients by using ridge regression only when the condition number is above a specific threshold (e.g., 10). The size of the ridge constant is chosen using an iterative searching procedure that finds the lowest positive number, *k*, which makes the recomputed condition number fall below the threshold.[4] By adding a small bias, the correlation between the regressors will decrease and so will the condition number.

4. Data and methodology

To illustrate our NS term structure fitting procedure, we use Euribor rates maturing from 1 week up to 12 months and Euro swap rates with maturities between 2 years and 10 years. The Euro Overnight Index Average (EONIA), and the 20-, 25- and 30-year Euro swap rates were also collected to assess the out-of-sample extrapolation quality of the selected estimation procedures. The Euribor and EONIA rates were obtained from their official website, and the Euro swap rates were gathered from the Thompson DataStream®. Our dataset spans the period from January 4, 1999 to May 12, 2009 and includes 2644 days.

We use the smoothed bootstrap to construct the spot rate curves:

1. The spot rates with maturities less than 1 year are Euribor rates.
2. Swap rates are par yields, so we bootstrap the zero rates. Denote $S(\tau)$ as the swap rate and $R(\tau)$ as the Euribor rate, time to maturity being τ. The following equation helps us to extract the spot rates from the swap rates:

$$R(\tau) = \left[\frac{1 - S(\tau)\sum_{j=1}^{\tau-1}\frac{1}{[1+R(j)]^j}}{1+S(\tau)}\right]^{-1/\tau} - 1. \tag{5}$$

   Here $\tau = 2, \ldots, 10$.
3. As the relationship in Eq. (1) only holds for continuously compounded rates, we need to convert the annualized spot rates to continuously compounded rates:

$$r(\tau) = \log[1 + R(\tau)]. \tag{6}$$

Table 1 summarizes the descriptive statistics for some of the time series of continuously compounded spot rates we use to fit the spot rate curve. The table shows that the volatility of the time series increases from 0.98% for the weekly rates to 1.03% for the 3-month rates, and then goes down to 0.70% for the 10-year spot rates. The average spot rate increases as time to maturity grows, from 3.19% for the one-week rates, to 4.53% for a 10-year maturity. Autocorrelation is high for rates of all maturities, from above 0.998 with a 5-day lag to above 0.916 with a 255-day lag.

Our dataset contains a lot of variations in the level, the slope and the curvature of the term structure. The short-term spot rates (1-week) vary from approximately 3.75% to 5% in 2008, and decrease to almost 0.7% in 2009, due to the financial crisis. The long-term spot rate is relatively stable, varying between 3.13% and almost 6%. The spot rate curve is most often upward sloping. Around 2006 there are humps in the spot rate curves, while at other times there are troughs. Between 2003 and 2005 the spot rate curve is flatter compared to the S-shaped curves in other periods.

[4] We start with $k = 0$ and we iteratively re-compute the condition number after increasing the ridge constant with 0.001. We stop iterating when the recomputed condition number is lower than the pre-specified threshold (i.e., 10).

Please cite this article as: Annaert, J., et al., Estimating the spot rate curve using the Nelson–Siegel model, *International Review of Economics and Finance* (2013), http://dx.doi.org/10.1016/j.iref.2013.01.005

Table 1
Descriptive statistics of some spot rates (rates are in percentage).

Maturity	Mean	Std. dev.	Min.	Max.	$\widehat{\rho}(5)$	$\widehat{\rho}(25)$	$\widehat{\rho}(255)$
1 week	3.186	0.982	0.710	5.240	0.998	0.972	0.926
1 month	3.238	0.995	0.866	5.257	0.999	0.975	0.926
3 months	3.334	1.030	1.307	5.431	0.999	0.983	0.941
6 months	3.377	1.025	1.501	5.449	0.999	0.984	0.945
12 months	3.453	1.014	1.656	5.451	0.999	0.983	0.942
2 years	3.568	0.911	1.724	5.435	0.998	0.974	0.923
5 years	4.029	0.762	2.599	5.614	0.998	0.971	0.916
10 years	4.530	0.696	3.134	5.957	0.998	0.975	0.934

Note: Spot rates are expressed in percentage with continuous compounding. The sample period runs from January 4, 1999 to May 12, 2009, totaling to 2644 days. The spot rates with maturities less than 1 year are Euribor rates, whereas those with a maturity of more than 1 year are bootstrapped from Euro swap rates. $\widehat{\rho}(n)$ is the n-day lag autocorrelation.

5. Empirical comparison of the estimation methods

For every day in our time series we estimate the NS model based on the proposed estimation methods. In order to compare them, we evaluate their ability (1) to fit the empirical term structures ('in-sample') and (2) to extrapolate the contemporaneous EONIA, 20-, 25- and 30-year Euro swap rates ('out-of-sample'). As an evaluation criterion, we use the mean absolute error (MAE) measured in basis points, since this gives us a good indication of the economic importance of the results. The estimation procedure which produces (over the available time series) the lowest MAEs 'wins' the rat race.

5.1. The time series of the estimated parameters

First, we discuss the results for the grid search, and subsequently we comment on the parameter estimates for the OLS approach where the shape parameter is fixed. Finally, we present the parameters for the grid search using the conditional ridge regression.

5.1.1. The grid search

Fig. 2 graphically represents the time series of β_0, β_1, β_2 and $(\beta_0+\beta_1)$ estimates, based on a grid search using OLS. At some points in time, the three coefficients are clearly quite erratic. Moreover, for the long term interest rate level β_0, some negative values are obtained, thereby clearly violating any economic intuition. The short end of the term structure, denoted by $(\beta_0+\beta_1)$, is always positive. This is explained by a highly negative correlation between the time series of β_0 and β_1 estimates. For August 21, 2007 e.g., the high β_0 coefficient (9.737) is accompanied with a low β_1 coefficient (−5.095), which leads to an estimate of the short rate of 4.642%.

In an attempt to understand the source of the erratic time series behavior, we looked at the histogram of the estimated shape parameters. The variation of the shape parameter estimates indicates that this parameter cannot be assumed constant over time. More than 55% of the estimated shape parameters are located within the range of 0 to 2, approximately 20% within the range of 2 to 4, and a little more than 19% within the range of 8 to 10. For the shape parameters within the range of 8 to 10, 472 out of 506 are estimated at the upper bound of the search interval i.e., at 10.[5] There are two explanations for the relatively large amount of the shape parameter estimates at the upper bound of the search interval: (1) the absence of a hump/trough in which case the curvature component can simply be dropped from the NS specification or (2) the presence of more than one hump/trough in which case a more flexible model such as Svensson (1994) might be a more appropriate model specification.

Grid search not only results in the erratic time series of factor estimates, the precision of the estimates is also very time varying. Fig. 3 redraws (as an example) the estimates of the long term spot rate (solid line) and its standard errors (dashed line). Whereas the standard errors are small at times, many periods of turbulence are shown in which the standard errors become 1% and more!

5.1.2. The OLS with fixed shape parameter

Fig. 4 plots the time series of the estimated parameters conditional on a fixed shape parameter of 1.37 (as in Diebold & Li, 2006) and 3 (as in Fabozzi et al., 2005). In contradiction to the time series pattern of the parameters produced by the grid search, the time series of the estimates are very smooth. Eyeballing Fig. 4, smoothness seems to be related to the choice of the shape parameter. The higher shape parameter ($\lambda = 3$) results in a less smoothed time series of the estimated β_0, β_1 and β_2 than the lower one ($\lambda = 1.37$).

Taking into account that the long-term interest rate level implied by the NS model is now always positive, and the short end does not show negative interest rate estimates either, fixing the shape parameter seems to alleviate the mentioned problems. However, the economic interpretation of the coefficients remains problematic. At the start of our series e.g., the long term rate is estimated as being 4.53% ($\lambda = 1.37$) and 5.81% ($\lambda = 3$) whereas the 30 year swap rate was 4.99%. So, which shape parameter should we consider to be the most appropriate?

[5] We performed a grid search with the shape parameter ranging from 0 to 10, 0 to 20 (not reported) or 0 to 30 (not reported). The shape parameters that are estimated to be at the upper bound (e.g., 10) are also estimated at the upper bound of the search interval when this upper bound is 20 or 30.

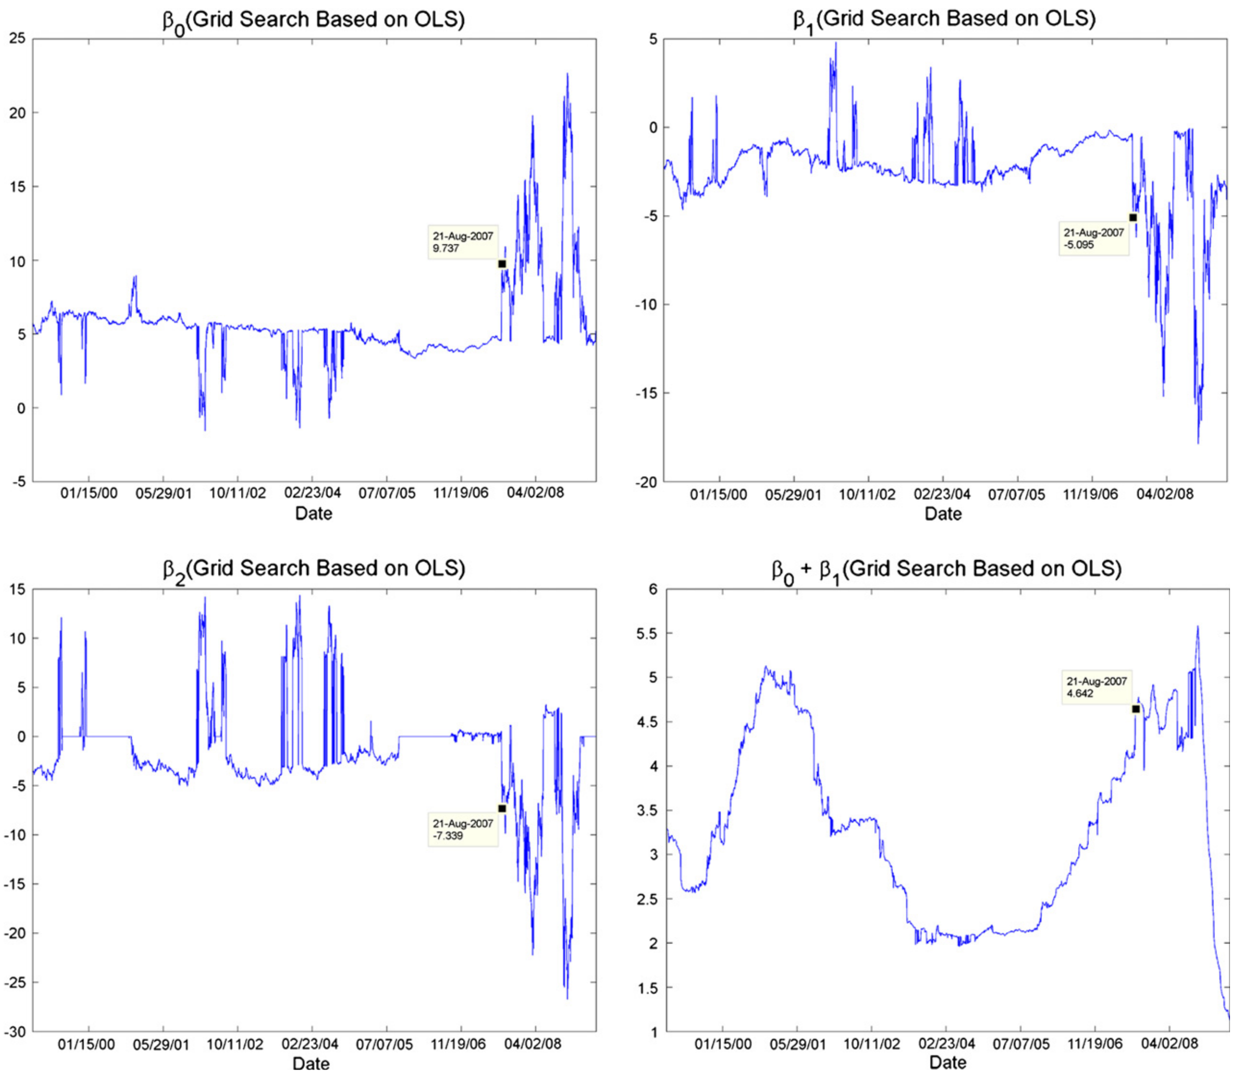

Note:This figure plots the time series of β_0, β_1, β_2 and ($\beta_0 + \beta_1$) estimates of the NS model on Euro spot rates based on a grid search using OLS over the period from January 4, 1999 to May 12, 2009.

Fig. 2. Time series of estimated parameters with the grid search based on OLS.

The precision of the estimates also differs dramatically over the various methods used. Fig. 5 shows the standard errors for the long term rate using $\lambda = 1.37$. In this case, the standard error went up to almost 40 basis points during the financial crisis whereas the grid search produced standard errors up to 700 basis points (see Fig. 3). For $\lambda = 3$, standard errors increased during the crisis up to 90 basis points. Whether the shape parameter can best be fixed to 1.37 or to 3, however, remains an open question. Our estimates suggest that the economic characteristics of the time series of the estimated coefficients may be quite different. Taking the variability of the shape parameter estimates we obtained in our grid search into account, it can be questioned whether the time variation in λ can be ignored!

5.1.3. Grid search with conditional ridge regression

A drawback of the ridge regression technique is the lack of standard errors, which prohibits any kind of significance tests on the estimated coefficients (DeMaris, 2004). However, we can (visually) examine the stability of the time series estimates. Fig. 6 panel A shows that a low fixed shape parameter does smooth the extreme jumps in the coefficient series almost completely, whereas the ridge regression takes a middle position. Whenever the condition number does not exceed the threshold, no multicollinearity problems occur and the ridge regression is redundant. Whenever the correlation between the regressors exceeds our threshold, the ridge regression has a 'moderate' smoothing effect. Ridge regression will not change the correlation between the regressors but the estimated coefficients β_0, β_1 and β_2 will be affected by the penalty function in Eq. (4), which will have a smoothing effect. In our application the implicit cut-off rate for the correlation used is −83.32%. Fig. 6 panel B plots the correlation between the two regressors

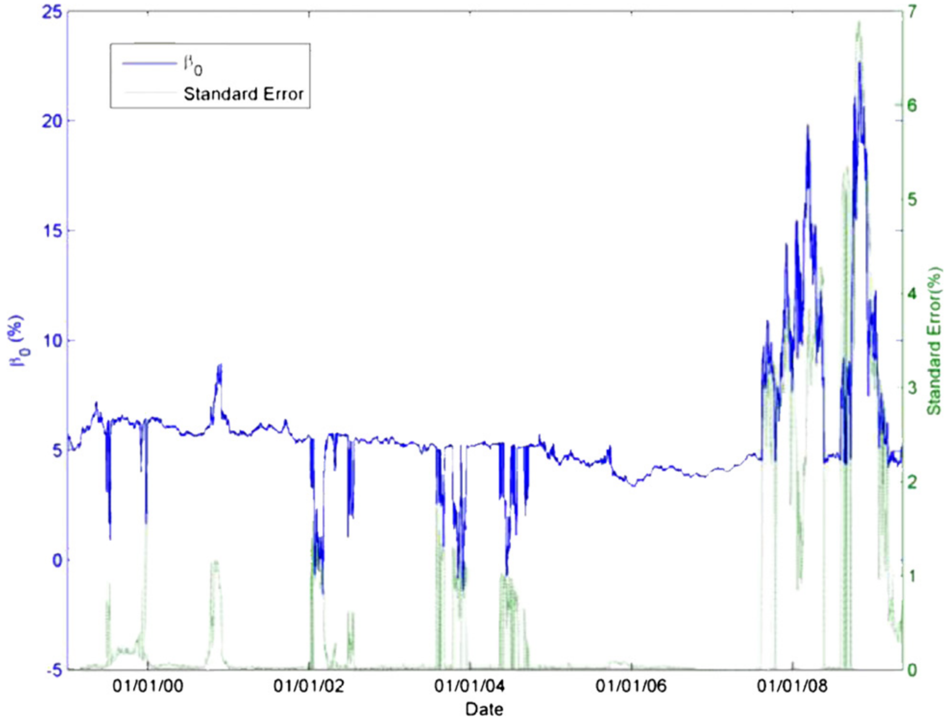

Note:This figure plots the long term spot rate estimates (solid line) and their standard errors (dashed line) based on grid search over the sample period from January 4, 1999 to May 12, 2009, totaling to 2644 days. The erratic behaviour of the NS model based on grid search is illustrated.

Fig. 3. Long term spot rate estimates and their standard errors — grid search.

based on the λ's determined by the grid search. It also plots the correlation used when λ = 3 i.e., −93.13%.[6] The figure shows that for a substantial amount of the days in our sample (i.e., 28.59%) the correlation in absolute terms is higher than 93.13%. Hence, for a lot of days the NS term structure has an inherent multicollinearity problem which the fixed λ procedures completely ignore. The estimated grid search based λ'*s* clearly show that the fixed λ assumption is simply too strict. By examining the extrapolated EONIA and long term swap rates, we show that fixing λ comes at a price (see Section 5.3).

Fig. 7 shows the time series of all the estimated coefficients using the ridge regression and the grid search based on OLS.

The estimates from the ridge regression are more stable compared to the results from the grid search. There are no negative values in the long-term interest rate level anymore. The long-term interest rate level jumped up from 4.80% in June 2007 to 5.85% in August 2007, then 8.14% in March 2008, and then reached its peak of 8.99% in October 2008. Afterwards the long-term interest rate level went down. The ridge regression improves the stability of the estimates calculated by the grid search. And the positive long-term interest rate level complies with the economic intuition behind the NS model. The short end of the term structure is again always positive, consistent with reality.

In order to measure the stability of the time series of estimated coefficients more formally, we compute the standard deviation of their first differences (Table 2). We notice that the ridge regressions have a lower volatility for all the three parameters. Moreover, an *F*-test at the 95% confidence interval shows that the standard deviations of the ridge regression coefficient changes are significantly lower than those obtained through the grid search. The ridge regression, hence, can substantially reduce the instability of the estimates in the grid search. Compared with the fixed shape parameter procedures, the ridge regression allows the shape parameter to vary over time.

5.2. In-sample performance

In order to examine the in-sample performance, we compute the mean absolute errors between the fitted and the bootstrapped spot rate (Table 3). To test whether the MAEs are statistically different from each other, we compute:

$$\widehat{\alpha}\cdot\mathbf{1} = \left|\mathbf{r}- \ \widehat{\mathbf{r}}_{Method1}\right| - \left|\mathbf{r}- \ \widehat{\mathbf{r}}_{Method2}\right|, \tag{7}$$

where $\mathbf{r}$ is the vector of empirical rates for a certain time to maturity, $\mathbf{1}$ is a vector of ones, $\widehat{r}_{Method1}$ and $\widehat{r}_{Method2}$ are the vectors of

[6] Section 3.1 reports that Diebold and Li (2006) force the correlation to be −54.94%.

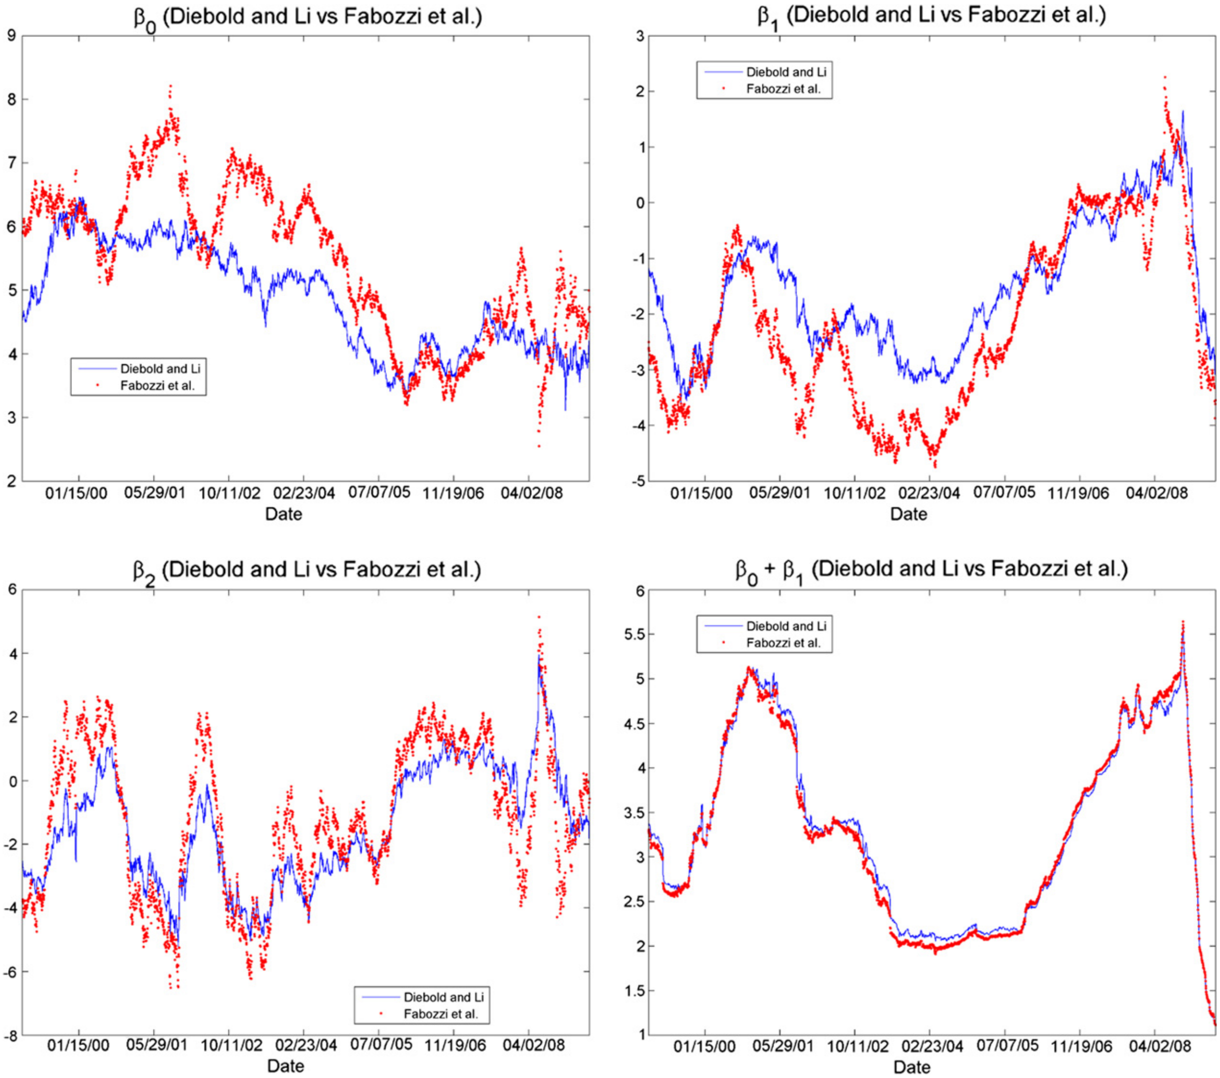

Note:This figure plots the time series of β_0, β_1, β_2 and ($\beta_0 + \beta_1$) estimates of the NS model on Euro spot rates based on fixed shape parameters of 1.37 (line) and 3 (dot). The sample period runs from January 4, 1999 to May 12, 2009.

Fig. 4. Time series of estimated parameters with fixed shape parameters.

estimated spot rates using different methods (e.g., grid, ridge, DL and FMP). If the estimated coefficient $\widehat{\alpha}$ is significantly negative, then we consider *Method 1* to be better than *Method 2*. The Newey–West correction is used to remove serial correlation from the residuals.

Although the overall MAE is minimized by the construction for the grid search, the in-sample MAEs per maturity show how the fixed shape parameter methods perform compared to the ridge regression technique. The MAE of FMP (DL) is for 13 (9) out of the 22 maturities the highest. For 8 out of the 22 maturities, the ridge regression generates the lowest MAE, but it never produces the highest MAEs. The grid search is for more than half of the maturities of the MAE minimizing method. The in-sample MAEs reveal the limitations of reducing the model flexibility by fixing the shape parameters. Ridge regression outperforms these two methods convincingly.

5.3. Out-of-sample performance

To investigate the ability of these four approaches to extrapolate the long and short ends of the term structure, we will compare the estimated rates to the EONIA, the 20-, 25- and 30-year Euro swap rates. The MAE will again be our criterion.

Theoretically speaking, the estimated $\beta_0 + \beta_1$ represents the short end of the term structure. However, the EONIA rates are the shortest-term spot rates that can be observed in the market. Consequently, if a model can extrapolate the EONIA rates with the highest accuracy, it will be superior in extrapolating the short end of the spot rate curve.

Please cite this article as: Annaert, J., et al., Estimating the spot rate curve using the Nelson–Siegel model, *International Review of Economics and Finance* (2013), http://dx.doi.org/10.1016/j.iref.2013.01.005

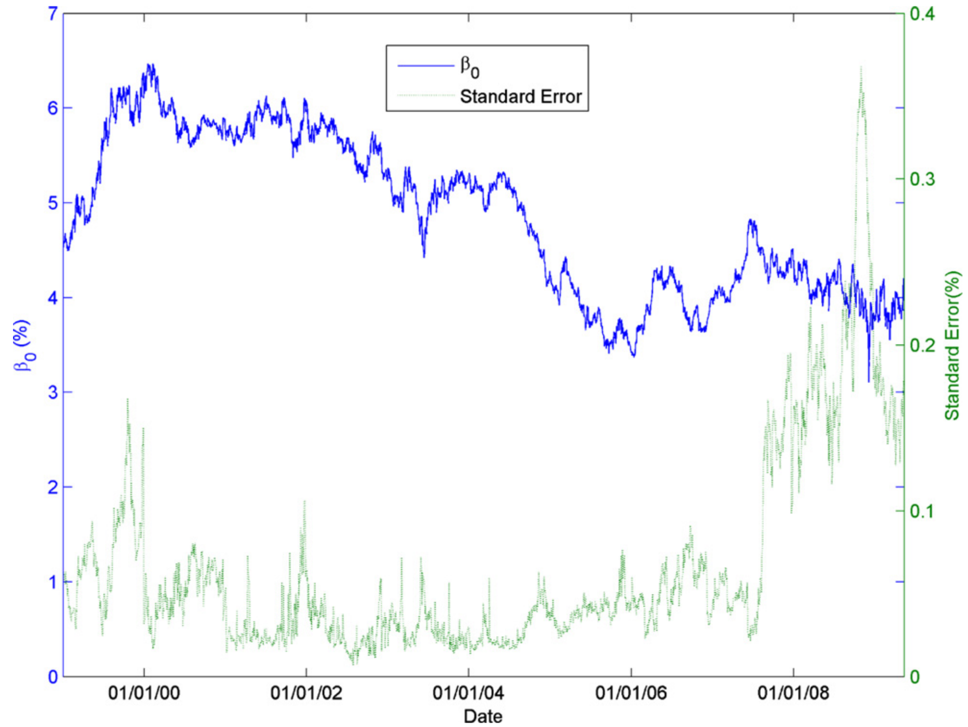

Note: This figure depicts the estimates of the long term rate (β_0) and their standard errors (Standard Error) based on a fixed shape parameter of 1.37 over the sample period from January 4, 1999 to May 12, 2009.

Fig. 5. Estimates of the long term rate and their standard errors ($\lambda = 1.37$).

The estimated EONIA rates are calculated by assuming a time to maturity of 1 day into Eq. (1) along with the estimated parameters for all four methodologies. Afterwards, the differences between the empirical and the estimated rates are examined following the same procedure as for the in-sample MAE tests.

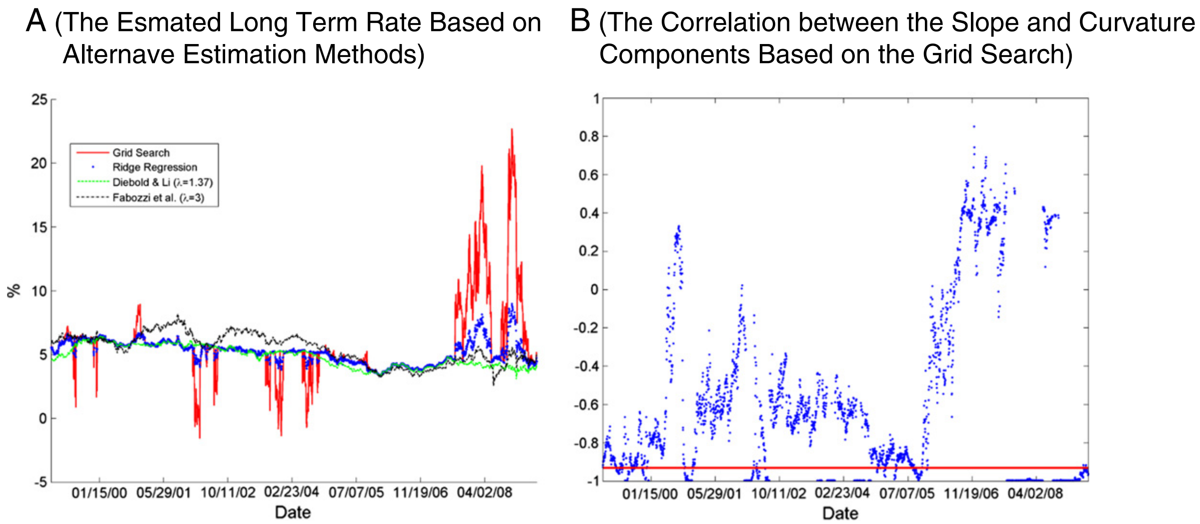

Note: Panel A plots the estimated long term rates based on the grid search, the ridge regression, fixed shape parameter of 1.37, and fixed shape parameter of 3. Panel B plots the time series of the correlation between the slope and curvature component based on λ's using the grid search.

Fig. 6. The estimated long term rate based and the time series of correlation.

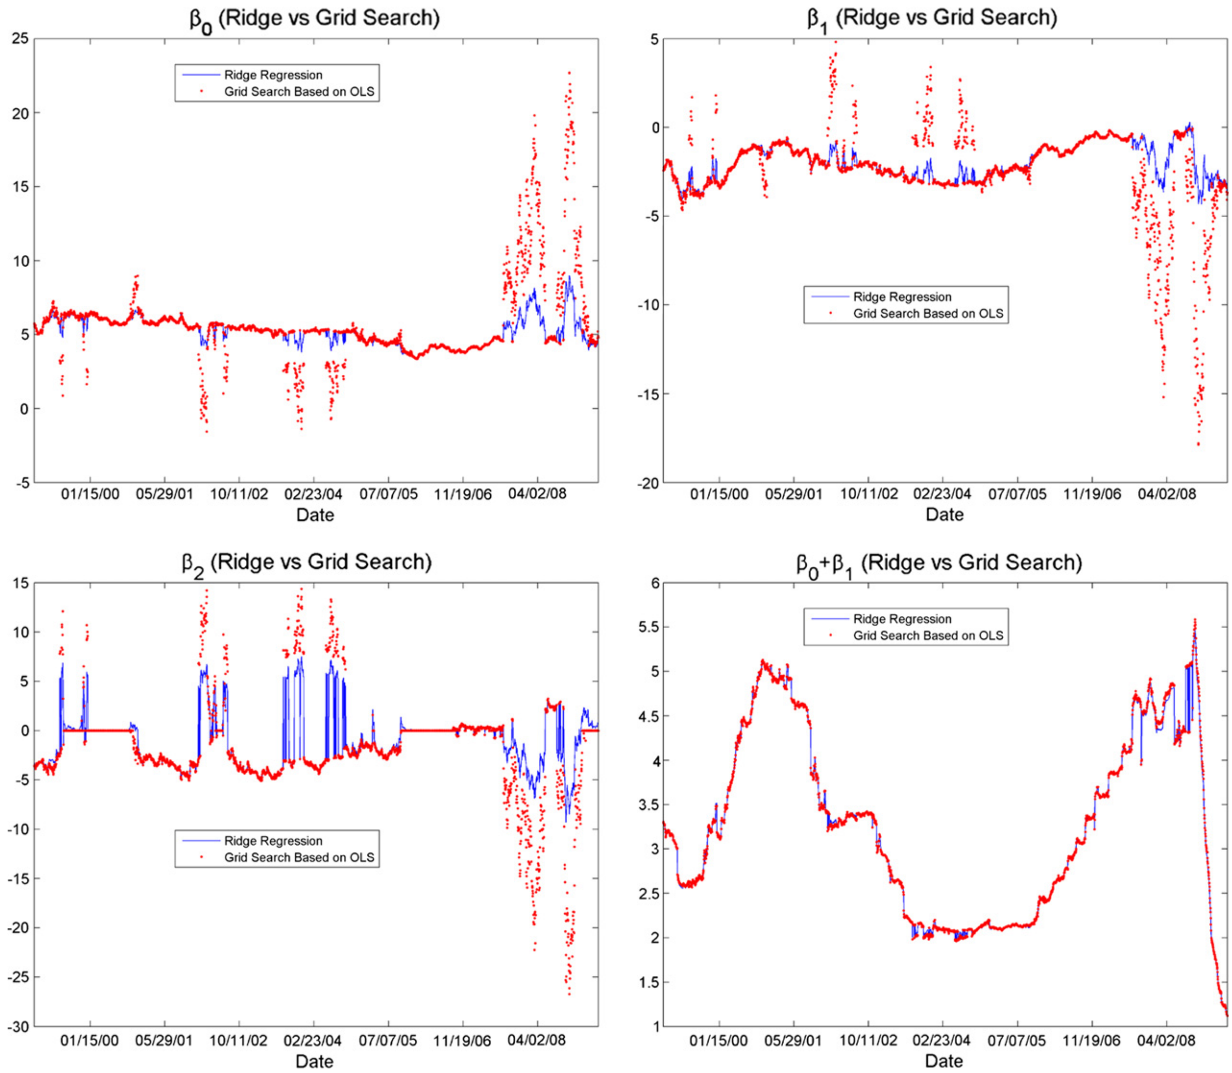

Note: This figure depicts the estimated parameters of the NS model based on both ridge regression (line) and grid search (dot). β_0 represents the long term interest rate level implied by the NS model. $\beta_0+\beta_1$ represents the short-term interest rate when time to maturity is zero.

Fig. 7. Estimated parameters with the ridge regression and grid search.

In addition, we also check which model can best fit the long-term end of the term structure. However, long-term spot rates are not directly observable in the market either. Moreover, due to the lack of swap rates for certain times to maturity (e.g., 11- or 29-year swap rate) to bootstrap the long-term spot rate (e.g., 30-year spot rate), we will use the following procedure to check the extrapolation ability: (1) combining Eqs. (1) and (5), we calculate the estimated swap rates using

$$\widehat{S}(\tau) = \frac{1-\left[1+\widehat{R}(\tau)\right]^{-\tau}}{\sum_{j=1}^{\tau}\frac{1}{\left[1+\widehat{R}(j)\right]^{j}}}, \tag{8}$$

where $\tau = 2, \ldots, 30$ are the times to maturity. $\widehat{R}(\tau)$, which equals $e^{\widehat{r}(\tau)}-1$ is the estimated τ-year annually compounded spot rate, $\widehat{r}(\tau)$ is the estimated continuously compounded spot rate, and $\widehat{S}(\tau)$ is the estimated swap rate. (2) We calculate and test the MAEs between the empirical and the estimated swap rates by using the same procedure as previously.

Table 2
The standard deviations of first differences in the estimates (in percentage).

	Grid	Ridge	DL ($\lambda=1.37$)	FMP ($\lambda=3$)
β_0	0.602	0.141	0.058	0.079
β_1	0.584	0.138	0.068	0.080
β_2	1.594	1.106	0.126	0.225

Note: This table shows the standard deviations of the first-order changes in the time series of the estimated parameters. The dataset used to estimate the parameters is composed by 1-week, 1- to 12-month, and 1- to 10-year spot rates. An *F*-test at a 95% confidence interval shows that all the standard deviations are significantly different from each other.

The results are summarized in Table 4. The ridge regression always yields the statistically lowest MAEs. The restriction DL and FMP put on the shape parameter makes them underperform compared to both the ridge and the OLS regressions. Moreover, the ridge regression has lower MAEs when extrapolating the long end, compared to the short end of the term structure.

The superiority of the ridge regression procedure is not only statistically significant. Looking at the MAEs for the 30-year swap rate, the MAEs are lowered with 7 to 15 basis points. On the short end, the gain is more moderate, up to 4 basis points.

Since the ridge regression adds a small bias to the OLS, we also consider another out-of-sample performance test in order to judge the appropriateness of the various estimation methods. If we denote $\varepsilon_\tau = r_\tau - \widehat{r}$ as the residual between the spot rate r_τ observed from the market the estimated spot rate by the NS model $\widehat{r}_\tau$ for a certain time to maturity τ, the bias is estimated as the average of the ε_τ's. The out-of-sample bias on each day is measured by averaging the residuals after fitting the EONIA, 20-, 25- and 30-year swap rates by using the four methods. Table 5 clearly shows that the bias introduced in the ridge regression does not affect the bias in the extrapolated spot rates in a material way compared to the other estimation methods.

5.4. Robustness check on the extrapolation ability

We consider two robustness checks to confirm the outperformance of the conditional ridge regression procedure we propose. First, we want to verify whether our results are mainly driven by the 2008–2009 financial crisis which is part of our dataset. Second, we have shown that the multicollinearity problem is severely affected by the choice of the maturity vector. We will examine whether our results are robust to the choice of a different maturity vector.

Table 3
In-sample mean absolute errors.

Maturity	Grid	Ridge	DL ($\lambda=1.37$)	FMP ($\lambda=3$)
1 week	0.113	0.103[a]	0.130	0.170[b]
1 month	0.067	0.059[a]	0.076	0.114[b]
2 months	0.035	0.031[a]	0.033	0.062[b]
3 months	0.031[a]	0.035	0.034	0.042[b]
4 months	0.030[a]	0.034	0.036[b]	0.034
5 months	0.031[a]	0.035	0.039[b]	0.036
6 months	0.033[a]	0.035	0.043	0.047[b]
7 months	0.033[a]	0.034	0.043	0.055[b]
8 months	0.034[a]	0.034	0.044	0.063[b]
9 months	0.035	0.035[a]	0.045	0.070[b]
10 months	0.036	0.035[a]	0.045	0.075[b]
11 months	0.038	0.037[a]	0.045	0.079[b]
12 months	0.041	0.039[a]	0.046	0.083[b]
2 years	0.055[a]	0.068	0.073[b]	0.057
3 years	0.048[a]	0.065	0.077[b]	0.055
4 years	0.033[a]	0.052	0.072[b]	0.065
5 years	0.028[a]	0.044	0.065[b]	0.061
6 years	0.024[a]	0.030	0.046	0.047[b]
7 years	0.015	0.013[a]	0.017	0.026[b]
8 years	0.013	0.018	0.022	0.008[a]
9 years	0.020[a]	0.037	0.054[b]	0.034
10 years	0.033[a]	0.055	0.083[b]	0.066

Note: In-sample mean absolute errors between the produced data from the NS and the empirical data are presented. The dataset used to estimate the parameters is composed by 1-week, 1- to 12-month, and 1- to 10-year spot rates.

[a] This approach yields the lowest MAE for this time to maturity.

[b] This approach yields the highest MAE for this time to maturity. Except the following pairs, all the MAEs are significantly different from each other (at 95% confidence interval) based on the *t*-test with Newey–West corrected standard errors: 2-month: grid and DL; 3-month: ridge and DL; 4-month: ridge and FMP; 5-month: ridge and FMP; 2-year: grid and FMP; 5-year: DL and FMP; 6-year: DL and FMP; 9-year: ridge and FMP. The mean absolute errors are tested by using Eq. (7).

Please cite this article as: Annaert, J., et al., Estimating the spot rate curve using the Nelson–Siegel model, *International Review of Economics and Finance* (2013), http://dx.doi.org/10.1016/j.iref.2013.01.005

Table 4
Out-of-sample mean absolute extrapolation errors.

Maturity	Grid	Ridge	DL ($\lambda=1.37$)	FMP ($\lambda=3$)
Overnight EONIA	0.254	0.246[a]	0.287[b]	0.278
20-year swap rate	0.201	0.142[a]	0.234[b]	0.221
25-year swap rate	0.262	0.150[a]	0.236	0.268[b]
30-year swap rate	0.304	0.152[a]	0.227	0.307[b]

Note: Out-of-sample MAEs between the produced data from the NS and the empirical data are presented. The dataset used to estimate the parameters is composed by 1-week, 1- to 12-month, and 1- to 10-year spot rates.
[a] This approach yields the lowest MAE for this time to maturity.
[b] This approach yields the highest MAE for this time to maturity. Except the following pairs, all the MAEs are significantly different from each other (at 95% confidence interval) based on the *t*-test with the Newey–West correction on standard errors: 20-year: grid and FMP, DL and FMP; 25-year: DL and FMP, grid and FMP; 30-year: grid and FMP. The mean absolute errors are tested by using Eq. (7).

5.4.1. The impact of the financial crisis

As seen in Fig. 7, the financial crisis has had a substantial impact on parameter estimation. We thus divide our dataset into two subsets, the pre-crisis period from January 4, 1999 to July 2, 2007 (2169 days), and the crisis period from July 3, 2007 to May 12, 2009 (475 days). The out-of-sample extrapolation power of all four approaches (grid, ridge, DL and FMP) is presented in Table 6.

A *t*-test between the pre-crisis and crisis period MAEs at 95% confidence level shows that during the financial crisis the extrapolation ability of all the four methods is significantly lowered. The extrapolation ability of the grid search drops dramatically for both the short and the long ends of the term structure, while for the other methodologies, the financial crisis has more impact on the short end than on the long end. Nevertheless, the ridge regression performs consistently in both periods, while the grid search clearly does not. During the financial crisis the grid search based extrapolations are worst for the long-end of the spot rate curve whereas in the pre-crisis period, the FMP model is worst for these maturities. The results shown in Table 6 again confirm our previous findings that the ridge regression superiorly extrapolates both ends of the spot rate curve.

5.4.2. A different maturity vector

In order to test the robustness of our results, we estimate the spot rate curves again by using only 1-week, 6-month and 1 to 10-year spot rates. The results, shown in Table 7, also confirm our previous findings that the ridge regression is superior in extrapolating both ends for the spot rate curve. However, unlike the other dataset, now the extrapolation of FMP on 25- and 30-year swap rates is not the worst, but the grid search is. Here the correlation between the slope and hump factors using the

Table 5
Out-of-sample bias.

	Grid	Ridge	DL ($\lambda=1.37$)	FMP ($\lambda=3$)
Mean	−0.070	4.00E−4	0.099	−0.089
Std. dev.	0.384	0.152	0.102	0.214
Minimum	−2.216	−0.925	−0.356	−0.820
Maximum	0.703	0.369	0.393	0.558

Note: The bias is expressed in percentage. The sample period runs from January 4, 1999 to May 12, 2009, totaling to 2644 days for which the out-of-sample estimation bias is measured. The out-of-sample bias on each day is measured by averaging the residuals after fitting the EONIA, 20-, 25- and 30-year swap rates by using the four methods. All the biases are significantly different from each other based on a *t*-test with the Newey–West corrected standard errors.

Table 6
Out-of-sample mean absolute extrapolation errors (pre-crisis vs crisis period).

Maturity	Grid		Ridge		DL ($\lambda=1.37$)		FMP ($\lambda=3$)	
	Pre-crisis	Crisis	Pre-crisis	Crisis	Pre-crisis	Crisis	Pre-crisis	Crisis
Overnight EONIA	0.161	0.676	0.161[a]	0.635[a]	0.195[b]	0.704	0.169	0.775[b]
20-year swap rate	0.142	0.470[b]	0.137[a]	0.165[a]	0.206	0.363	0.217[b]	0.240
25-year swap rate	0.161	0.721[b]	0.140[a]	0.198[a]	0.218	0.317	0.260[b]	0.306
30-year swap rate	0.169	0.919[b]	0.131[a]	0.251[a]	0.211	0.296	0.296[b]	0.361

Note: Out-of-sample MAEs between the produced data from the NS and the empirical data are presented for the pre-crisis period from January 4, 1999 to July 2, 2007 (2169 days) and the crisis period from July 3, 2007 to August 12, 2009 (475 days). The dataset used to estimate the parameters is composed by 1-week, 1–12 month, and 1- to 10-year spot rates.
[a] This approach yields the lowest MAE for this time to maturity.
[b] This approach yields the highest MAE for this time to maturity. Pre-crisis: Except the following pairs, all the MAEs are significantly different from each other based on the *t*-test with the Newey–West correction on standard errors: EONIA: grid and ridge, grid and FMP, ridge and FMP; 20-year: grid and ridge, DL and FMP. Crisis: Except the following pairs, all the MAEs are significantly different from each other based on the *t*-test with the Newey–West correction on standard errors: EONIA: grid and DL; 20-year: grid and DL; 25-year: DL and FMP; 30-year: ridge and DL, DL and FMP.

Table 7
Out-of-sample mean absolute errors.

Maturity	Grid	Ridge	DL ($\lambda=1.37$)	FMP ($\lambda=3$)
Overnight EONIA	0.293	0.217[a]	0.255[b]	0.243
20-year swap rate	0.171	0.128[a]	0.205[b]	0.179
25-year swap rate	0.226[b]	0.140[a]	0.207	0.222
30-year swap rate	0.266[b]	0.147[a]	0.200	0.259

Note: Out-of-sample MAEs between the produced data from the NS and the empirical data are presented. The dataset used to estimate the parameters is composed by 1-week, 6-month, and 1- to 10-year spot rates.
[a] This approach yields the lowest MAE for this time to maturity.
[b] This approach yields the highest MAE for this time to maturity. Except the following pairs, all the MAEs are significantly different from each other based on the *t*-test with the Newey–West correction on standard errors: EONIA: grid and DL; 20-year: grid and FMP; 25-year: grid and DL, grid and FMP, DL and FMP; 30-year: grid and DL.

DL-estimation recipe is −0.35 instead of −0.55. The highest MAEs for the 25- and 30-year swap rates obtained using the grid search based on the OLS regression, can be blamed on the unstable estimates.

Again, we test the influence of the financial crisis on the performance of the proposed methods, this time on our limited sample. The results from this analysis comparing estimates of the pre-crisis and crisis periods are summarized in Table 8.

At the 95% confidence level, a *t*-test between the pre-crisis and crisis period MAEs shows that the performance of the four methods behaves similar to that of the other dataset. The high volatility of the estimates in the grid search makes its performance drop substantially during the financial crisis. The ridge regression always has the highest extrapolation ability except for the 30-year swap rate during the crisis period, where the DL has the highest extrapolation ability. However, the difference between these two methods for the 30-year swap rate during the financial crisis is not statistically significant.

6. Conclusion

Many researchers have reported problems in estimating the NS model. We have shown that multicollinearity between the slope and the hump factor is causing the instability of the regression coefficients over time as well as the large standard errors on the coefficients (in the case of the grid search) and extremely smoothed time series of the parameters (in the cases where the shape parameter is fixed). To alleviate these estimation problems, we apply the ridge regression technique, whenever the grid search based estimate of the shape parameter results in highly correlated slope and hump factors. For the Euro spot rate curves, over the period 1999–2009, we compare the grid search estimates – originally proposed by Nelson and Siegel – to our ridge regression estimates. The Diebold and Li (2006) and the Fabozzi et al. (2005) estimates were also calculated as a benchmark.

The in-sample comparison shows that the grid search produces erratic time series estimates which sometimes violate the economic intuition behind the NS model. The distribution of the freely estimated shape parameters and the in-sample fitting errors reveals the limitation of the use of a fixed shape parameter. The out-of-sample extrapolation at the two ends of the term structure shows that the ridge regression always produces the lowest mean absolute errors. For the long end of the term structure, the economic gain in the MAE mounted up to 15 basis points. For the short end, the differences in MAEs were statistically significant but economically smaller.

The robustness checks show that the ridge regression performs robustly and consistently better in different economic environments (pre-crisis and crisis periods) and with different choices of the maturity vector (i.e., the set of financial instruments used to bootstrap the spot rate curve).

Based on our findings, fixing the shape parameter in order to avoid multicollinearity, is a statistical trick that does reduce the correlation between the regressors when the fixing is judicially chosen. It however ignores the bare fact that in practice, the term

Table 8
Out-of-sample mean absolute extrapolation errors (pre-crisis vs crisis period).

Maturity	Grid		Ridge		DL ($\lambda=1.37$)		FMP ($\lambda=3$)	
	Pre-crisis	Crisis	Pre-crisis	Crisis	Pre-crisis	Crisis	Pre-crisis	Crisis
Overnight EONIA	0.216[b]	0.642[b]	0.156[a]	0.499[a]	0.186	0.571	0.160	0.624
20-year swap rate	0.128	0.365[b]	0.120[a]	0.164[a]	0.191[b]	0.268	0.173	0.205
25-year swap rate	0.150	0.578[b]	0.124[a]	0.213[a]	0.201	0.236	0.209[b]	0.280
30-year swap rate	0.161	0.749[b]	0.121[a]	0.264	0.192	0.236[a]	0.242[b]	0.337

Note: Out-of-sample MAEs between the produced data from the NS and the empirical data are presented for the pre-crisis period from January 4, 1999 to July 2, 2007 (2169 days) and the crisis period from July 3, 2007 to August 12, 2009 (475 days). The dataset used to estimate the parameters is composed by 1-week, 6-month, and 1- to 10-year spot rates.
[a] This approach yields the lowest MAE for this time to maturity.
[b] This approach yields the highest MAE for this time to maturity. Pre-crisis: Except the following pair, all the MAEs are significantly different from each other based on the *t*-test with the Newey–West correction on standard errors: 25-year: DL and FMP. Crisis: Except the following pairs, all the MAEs are significantly different from each other based on the *t*-test with the Newey–West correction on standard errors: EONIA: grid and DL, grid and FMP; 25-year: ridge and DL, DL and FMP; 30-year: ridge and FMP.

Please cite this article as: Annaert, J., et al., Estimating the spot rate curve using the Nelson–Siegel model, *International Review of Economics and Finance* (2013), http://dx.doi.org/10.1016/j.iref.2013.01.005

structures do take all kinds of (humpy) shapes and that the hump simply is not fixed over the time to maturity spectrum. The loss in flexibility comes at a price especially for the extrapolation of long term spot rates.

References

Bank of International Settlements (2005). Zero-coupon yield curves — Technical documentation. *BIS Paper*, *25*. (pp. 1–37).

Barrett, W. R., Gosnell, T. F., Jr., & Heuson, A. J. (1995). Yield curve shifts and the selection of immunization strategies. *Journal of Fixed Income*, *5*(2), 53–64.

Belsley, D. A. (1991). *Conditioning diagnostics — Collinearity and weak data in regression.* Wiley series in probability and mathematical statistics. New Jersey: Wiley-Interscience Publication.

Bliss, R., & Fama, E. (1987). The information in long-maturity forward rates. *The American Economic Review*, *77*, 680–692.

Cairns, A. J. G., & Pritchard, D. J. (2001). Stability of descriptive models for the term structure of interest rates with applications to German market data. *British Actuarial Journal*, *7*, 467–507.

Coroneo, L., Nyholm, K., & Vivada-Koleva, R. (2011). How arbitrage-free is the Nelson–Siegel model? *Journal of Empirical Finance*, *18*, 393–407.

de Pooter, M. (2007). Examining the Nelson–Siegel class of term structure models. *Tinbergen Institute Discussion Paper IT, 2007-043/4* (pp. 1–56).

DeMaris, A. (2004). Regression with social data: Modeling continuous and limited response variables*Wiley series in probability and statistics*. New Jersey: Wiley-Interscience Publication.

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

Diebold, F. X., Li, C., & Yue, V. Z. (2008). Global yield curve dynamics and interactions: A generalized Nelson–Siegel approach. *Journal of Econometrics*, *146*, 351–363.

Dullmann, K., & Uhrig-Homburg, M. (2000). Profiting risk structure of interest rates: An empirical analysis for Deutschemark-denominated bonds. *European Financial Management*, *6*, 367–388.

Central Bank, European (2008). The new Euro area yield curves. *Monthly Bulletin* (pp. 95–103).

Fabozzi, F. J., Martellini, L., & Priaulet, P. (2005). Predictability in the shape of the term structure of interest rates. *Journal of Fixed Income*, *15*(1), 40–53.

Ferguson, R., & Raymar, S. (1998). A comparative analysis of several popular term structure estimation models. *Journal of Fixed Income*, *7*(4), 17–33.

Gurkaynak, R. S., Sack, B., & Wright, J. H. (2006). The U.S. Treasury yield curve: 1961 to the present. *Finance and Economics Discussion Series, 2006–28*. (pp. 1–41).

Haile, F., & Pozo, S. (2008). Currency crisis contagion and the identification of transmission channels. *International Review of Economics and Finance*, *17*, 572–588.

Hodges, S. D., & Parekh, N. (2006). Term structure slope risk: Convexity revisited. *Journal of Fixed Income*, *16*(3), 54–59.

Hurn, A. S., Lindsay, K. A., & Pavlov, V. (2005). Smooth estimation of yield curves by Laguerre function. Paper presented at international congress on modelling and simulation advances and applications for management and decision making. http://eprints.qut.edu.au/23852/1/23852_hurn_2006001493.pdf.

Kutner, M. H., Nachtsheim, C. J., Neter, J., & Li, W. (2004). *Applied linear statistical models* (5th ed.). New York: McGraw-Hill.

Lee, H. W., Xie, Y. A., & Yau, J. (2011). The impact of sovereign risk on bond duration: Evidence from Asian sovereign bond markets. *International Review of Economics and Finance*, *20*, 441–451.

Litterman, R., & Scheinkman, J. (1991). Common factors affecting bond returns. *Journal of Fixed Income*, *1*(1), 54–61.

Martellini, L., & Meyfredi, J. C. (2007). A copula approach to value-at-risk estimation for fixed-income portfolios. *Journal of Fixed Income*, *17*(1), 5–15.

McCulloch, J. H. (1971). Measuring the term structure of interest rates. *Journal of Business*, *44*, 19–31.

McCulloch, J. H. (1975). The tax-adjusted yield curve. *Journal of Finance*, *30*, 811–830.

Nelson, C., & Siegel, A. F. (1987). Parsimonious modeling of yield curves. *Journal of Business*, *60*, 473–489.

Seber, G. A. F., & Wild, C. J. (2003). *Nonlinear regression.* Wiley series in probability and statistics. New Jersey: Wiley-Interscience Publication.

Shea, G. (1984). Pitfalls in smoothing interest rate term structure data: Equilibrium models of spline approximations. *Journal of Financial and Quantitative Studies*, *19*, 253–269.

Svensson, L. E. O. (1994). *Estimating and interpreting forward interest rates: Sweden 1992–1994.* IMF working paper, WP/94/114.

Svensson, L. E. O. (1996). Estimating the term structure of interest rates for monetary policy analysis. *The Scandinavian Journal of Economics*, *98*, 163–183.

Vasicek, O. A., & Fong, H. G. (1982). Term structure modeling using exponential splines. *Journal of Finance*, *73*, 339–348.

Vermani, V. (2012). On estimability of parsimonious term structure models: An experiment with the Nelson–Siegel specification. *Applied Economics Letters*, *19*, 1703–1706.

Yan, A. C., Shi, J., & Wu, C. (2008). Do macroeconomic variables matter for pricing default risk? *International Review of Economics and Finance*, *17*, 279–291.

Please cite this article as: Annaert, J., et al., Estimating the spot rate curve using the Nelson–Siegel model, *International Review of Economics and Finance* (2013), http://dx.doi.org/10.1016/j.iref.2013.01.005