Calibrating the Dynamic Nelson-Siegel Model: A Practitioner Approach

Ibanez, Francisco

Central Bank of Chile

14 December 2015

Calibrating the Dynamic Nelson-Siegel Model: A Practitioner Approach

Francisco Ibáñez
Central Bank of Chile

December 14, 2015

Abstract

The dynamic version of the Nelson-Siegel model has shown useful applications in the investment management industry. These applications go from forecasting the yield curve to portfolio risk management. Because of the complexity in the estimation of the parameters, some practitioners are unable to benefit from the uses of this model. In this note we present two approximations to estimate the time series of the model's factors. The first one has a more technical aim, focusing on the construction of a representative base to work, and uses a genetic algorithm to face the optimization problem. The second approximation has a practitioner spirit, focusing on the easiness of implementation. The results show that both approximations have good fitting for the U.S. Treasury bonds market.

1 Introduction

Over the last decades, modelling the yield curve has become a major interest for both academics and investment practitioners. Jones (1991) indicates that approximately 95% of the total return of a US Treasury bonds portfolio can be explained by three type of movements. These movements include parallel shifts (86%), slope twists (9.8%) and butterfly type movements (3.6%). Litterman and Sheinkman (1991) find similar results and expose the benefits of a three-factor approach (level, slope, curvature) on bond portfolio hedging.

There are many techniques and models to fit a yield curve, but very few of them end up being as popular as the model proposed by Nelson and Siegel (1987). Diebold and Rudebusch (2013) attribute this success to three main features. Firstly, the model respects the restrictions imposed by the economic and financial theory (e.g. discount factor tends to zero as maturity grows). Secondly, its parsimonious approximation avoids in-sample overfitting, improving its forecasting capacity. Finally, it can take any yield curve form, empirically observed in the market.

Thanks to the great versatility shown by the Nelson-Siegel model and the recent interest of researchers, many extensions of the model have emerged. Some of them include additional variables (Svensson, 1995), and no-arbitrage conditions (Christensen, Diebold and Rudebusch, 2011). Diebold, Ji and Li, (2006) extend the use of the model and show that managing interest rate risk with a three-factor approach (level, slope and curvature durations), can be more reliable than a one-factor approach (modified duration and convexity).

Although this model could be a powerful tool in the hands of portfolio managers, estimation techniques can result somewhat unfamiliar and technical. This study is organized as follows: the second section introduces, briefly, the Nelson-Siegel model. The third section develops a model calibration for the US market. The fourth section proposes a practitioner version of the exercise, more pragmatic and replicable. The fifth section concludes.

2 Modelling the yield curve

Nelson and Siegel (1987) modelled the yield curve using three components. The first one remains constant when the term to maturity (τ) varies. The second factor has more impact on short maturities. The impact of the third factor increases with maturity, reaches a peak and then decays to zero. The authors conclude that each one of these three factors governs the part of the yield curve where they have more impact; long (β_1), short (β_2) and medium term (β_3), respectively. There is a fourth factor (not entirely described), λ, that can be interpreted as a decay factor. The value of λ affects the fitting power of the model for different segments of the curve.

$$y(\tau) = \beta_1 + (\beta_2 + \beta_3)\left(\frac{1 - e^{-\lambda\tau}}{\lambda\tau}\right) - \beta_3 e^{-\lambda\tau} \tag{1}$$

Diebold and Li (2006) propose a dynamic version of the model. Firstly, they reinterpret the three components of the original model as level (l_t), slope (s_t) and curvature (c_t), at period t. Secondly, they realise that if the yield curve evolves in time, so do the factors explaining it. This allows the potential use of the model for forecasting purposes.

$$y_t(\tau) = l_t + s_t\left(\frac{1 - e^{-\lambda_t\tau}}{\lambda_t\tau}\right) + c_t\left(\frac{1 - e^{-\lambda_t\tau}}{\lambda_t\tau} - e^{-\lambda_t\tau}\right) \tag{2}$$

λ introduces a non-linearity in the equation (2). Nelson and Siegel (1987) note that by fixing this parameter, the problem assumes a lineal form, and it can be solved by OLS. Following this, Diebold and Li (2006) propose the two-step estimation. The first step of the procedure consists of fixing the decay factor to estimate the parameters. The second step consists of fitting a dynamic model to the factors $\left\{\hat{l}_t, \hat{s}_t, \hat{c}_t\right\}$ with a VAR or other relevant model. Diebold and Rudebusch (2013) analyse various methodologies to estimate the parameters of the model. The authors note that although more complex methodologies are superior, in principle, they have only slightly better fitting results than the two-step estimation.

This study has the objective of providing tools to estimate the first part of the two-step methodology. In other words, we intend to fix λ, so we can estimate the time series of level, slope and curvature. We will present two approaches. The first one will be a more rigorous approach, from the academic point of view. The second one will be more pragmatic, focusing on the easiness of implementation. Both methodologies will focus on the US Treasury market, due to its high liquidity and data availability.

3 Rigorous approach

This approach works on a rich price base, an ideal scenario to calibrate the model. Though its representation of the bond market is the best and its results are very good, the construction of the price base is time consuming and the estimation is computationally demanding. The rigorous methodology is described as follows:

- Using a secondary market price base of bills, notes and bonds issued by the US Treasury, over a determined window, we construct a panel of yield curves.
- Bootstrapping the yield curve of each period, we find the spot rates.
- We fit the model on each observation, without fixing the decay factor. As a result, we get a distribution of λ over the estimation window.
- Finally, we choose the one λ that shows the best fitting power over the whole window. Then, we can construct the time series for the factors of level, slope and curvature.

3.1 Data

Choosing the securities to be included in the estimation is no trivial task. Gürkaynak, Sack and Wright (2007) estimate the parameters of the Nelson-Siegel-Svensson model for the US Treasury yield curve from 1961 to 2006[1]. Although the authors intend to work on a rich sample that, ideally, includes every available security at every observation, they recognise that not all Treasuries are comparable in terms of liquidity, and specific features (e.g. callable bonds). To address this issue they set a few rules to filter these undesirable securities. Even though our work has a different aim than the cited study (the authors focus on a cross-sectional fit), the treatment of the data under this approach will be similar.

We set a daily frequency window between December 1999 and December 2013. This gives us 3,523 observations. For each period we gather the prices and features of the securities that had been part of the Barclays U.S. Treasury

[1]The resulting yields and the estimated parameters are being constantly updated, and can be accessed on the website http://www.federalreserve.gov/econresdata/feds/2006/.

Figure 1: Spot yield surface

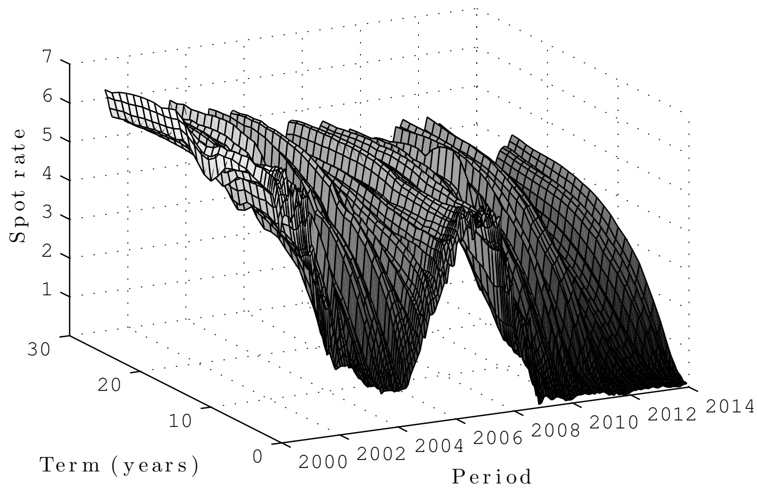

After bootstrapping each yield curve in the estimation window, we interpolate constant maturities from 1/4 to 25 years. Each one of these spot rates are comparable over time.

Index and the Barclays U.S. Treasury Bill Index during that day[2]. The first includes notes and bonds issued by the U.S. Government, excluding TIPS and STRIPS. The second index contains bills with less than a year of maturity. The less populated observation has 112 bonds and bills, while the most populated has 268. On average, we have 166 securities in each yield curve, making this an almost ideal price base to work with.

For each observation we bootstrap the yield curve, to obtain the spot rates. The roll-down effect and new issuances make it impossible to compare each curve directly. To face this problem, Diebold and Li (2006) interpolate the terms to maturity 3, 6, 9, 12, 15, 18, 21, 24, 30, 36, 48, 60, 72, 84, 96, 108 and 120 months. Using cubic splines, we extend their work adding maturities up to 25 years, spaced 12 months from each other. We do not reach the 30 years because there are some periods when the Treasury stopped issuing 30-year bonds. Additionally, and following market practices, we express term to maturity in years. As a result we obtain a spot rate surface with 32 constant maturities, shown in figure 1.

3.2 Estimation

Our next objective is to find a fixed value of λ that provides the best fit for every observation along the window. On each observation we seek to minimize

[2]Using both indexes we leave the Treasury notes with maturity less than a year out of the sample. Gürkaynak et al. (2007) argue that short maturity notes often behave oddly, probably due to the lack of liquidity, and segmented demand.

Figure 2: Correlation between level and slope loadings

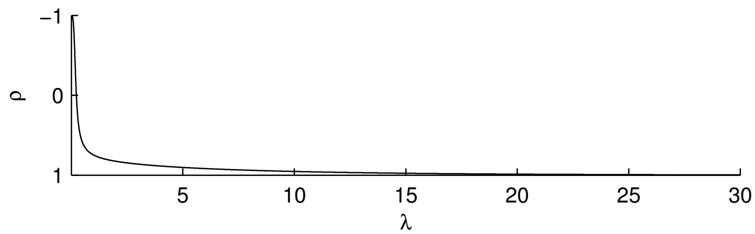

We graphed the correlation between the level and slope loadings (vertical axis), for each value of λ_t. The figure shows that the correlation tends to -1 when the decay factor tends to 0. On the other hand, the correlation tends to be perfect for λ_t over 30.

the root mean squared error (RMSE) between the yields generated by the model and the spot rates bootstrapped on the previous section. Although this thinking appears to be straight forward, it is necessary to impose some restrictions on the excercise. By doing this, we can improve the economic interpretation of the factors and use less resources solving the optimization problem.

$$\lim_{\tau \to \infty} y_t\left(\tau\right) = l_t$$

While the term to maturity (τ) increases, the equation (2) tends to the level factor. According to data published by the United States Treasury, the highest level recorded for a 30-year bond is 15.2% (September 1981), and has not exceeded levels over 10% since October 1985. Consistenly with this, we cap l_t at 15. Negative values for the level factor will imply a negative long term rate, which is inconsistent with the financial theory. For this reason, we set the lower limit to zero. By construction, this restriction affects the values that s_t and c_t can assume. Both can take positive and negative values, so we set their limits at -15 and 15[3].

A higher correlation between the slope and curvature factor loadings will affect the identification of the factors. Gilli, Grobe and Schumann (2010) state that a lower identification will translate in factors exchanging their values over time. Figure 2 shows that correlation between the level and slope loadings tends to be perfectly inverse when λ_t is close to zero. The correlation rapidly increases with the decay factor, becoming near perfect for values over 30. In order to address the collinearity and identification problem, we will set limits for λ_t between 0 and 30.

$$\lim_{\tau \to 0} y_t\left(\tau\right) = l_t + s_t$$

[3]If the instantaneous rate ($y_t(0)$) is zero, the maximum value that the slope ($y_t(\infty) - y_t(0)$) can assume is the level factor. Analogously, if the instantaneous rate is 15, the slope will be -15, thanks to the restriction of positive long term rate. Similar logic can be applied to set the limit of the curvature, so they will share the same limits.

When the term to maturity tends to zero (instantaneous interest rate), the yield provided by equation (2) approximates to the sum of the level and slope factors. With the objective of having a positive overnight rate, the sum of both factors must be positive.

Grouping all restrictions, we are going to resolve the following optimization problem:

$$\min_{\theta} Z = \frac{1}{\sqrt{m}} \|y_t - \hat{y}_t\|_2 \tag{3}$$

$$\theta = \{l_t, s_t, c_t, \lambda_t\}$$

$$s.t. \begin{cases} 0 < l_t \leq 15 \\ -15 \leq s_t \leq 15 \\ -15 \leq c_t \leq 15 \\ 0 < \lambda_t \leq 30 \\ 0 < l_t + s_t \end{cases}$$

where y_t is a vector containing the observed spot yields at period t, while $\hat{y}_t$ is a vector that contains yields provided by the equation (2), for each of the m terms of the curve.

Gilli et al. (2010) identify three major problems with estimating this model. Firstly, the optimization problem is nonconvex, showing multiple local optima. Secondly, the model is badly-conditioned for a certain range of parameters, thus the estimates are unstable before small price movements. Finally, the value of the decay factor directly affects the correlation between the loadings of the slope and curvature factors. This could be a problem for the forecasting use of the model. Additionally, if the model was correctly specified, the parameters would have a clear economic interpretation, making restrictions unnecessary.

The authors expose the benefits of an evolutionary algorithm to address the nonconvex optimization. Following this, we use a stochastic optimization based on a genetic algorithm. This method takes inspiration from the evolutionary process. The algorithm starts with a given initial population that evolves with every iteration (generation). It can do this by elitism (natural selection), mutation, and cross-over. This way, the algorithm has the capacity to face complex optimization problems and approach a global optima more easily. In order to reduce the average number of iterations and make the exercise less computationally demanding, we will use a hybrid approach. This technique uses the stochastic optimization to get rid of local optima, and applies a gradient based optimization. For this problem, we will use the SQP method, after the genetic algorithm.

For each one of the 3523 observations we estimate the set $\{l_t, s_t, c_t, \lambda_t\}$that best fit market spot rates. Figure 3 shows a histogram with optimal values for

Figure 3: Optimal values for Lambda

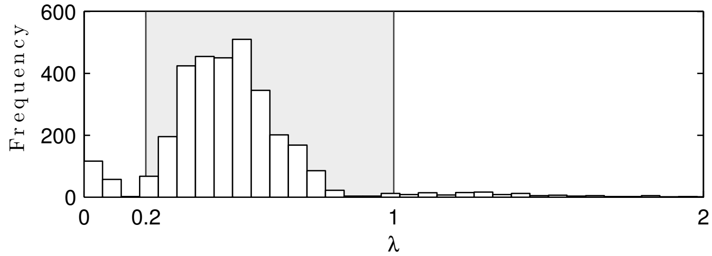

The histogram shows the optimal values for λ_t over the window. We define a search range between 0.2 and 1.

λ_t, along the estimation window. Analizing the results, we define a search space between 0.2 and 1, for two reasons. Firstly, the left region is too close to the minimum limit for λ_t. Secondly, and more importantly, these values give a high negative correlation between the factor loadings of slope and curvature. Alfaro, Becerra and Sagner (2011) use a search grid to define the decay factor. On each point of the grid, they fix λ_t at that value and estimate the other parameters linearly. Then, they choose the fixed value of lambda that gives the best fitting in terms of the RMSE, BIC, AIC and R^2. We follow a similar procedure over our search space, but choosing the value for the decay factor that gives the lowest mean absolute error (MAE), over the entire sample.

In a daily window between December of 1999 and December of 2013, the value for λ_t that produces the best fit, in terms of MAE, is 0.4835. This value generates close to 8.97 basis points of mean absolute deviation from the market rates, and it can be seen in Figure 4. If we follow Diebold and Li (2006) and fix the value of the decay factor at 0.7308[4], we get 12.12 basis points of error. We recognize two possible reasons for this. Firstly, the authors do not perform an optimization excercise to find this value. Instead, they fix λ_t at the value that maximizes the loading on the curvature factor at exactly 30 months. Secondly, their excercise considers maturities up to 10 years. Because we fit the model on rates up to 25 years, there could be a loss of fit on the longer maturities. Because the methodology proposed on this study is able to deliver better results than using the decay factor of the original work, we are content with it.

4 Practitioner Approach

Although the methodology described in the previous section shows a strong goodness of fit, the amount of time required to construct the base to work with is high, and the optimization procedure can become computationally demanding. Now we present a more pragmatical methodology that can be replicated by

[4]The authors propose 0.0609. Because they measure term to maturity in terms of months, we annualize their value.

Figure 4: Mean absolute error

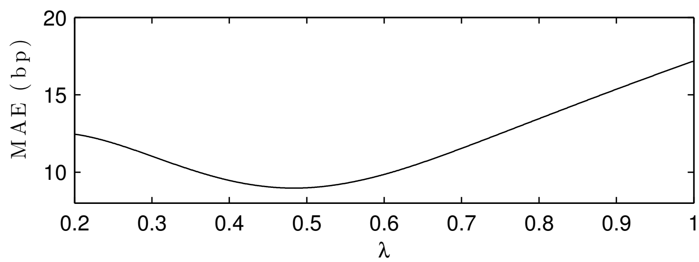

investment analysts. The practitioner methodology is described as follows:

- Using public and free data, we construct a spot rates surface.
- Then, we find a fixed value for λ, without a resource consuming calibration. This way, the methodology remains as simple as possible.
- Finally, we estimate the time series for the factors of level, slope and curvature. We use the restricted least squares (RLS) method, so we can incorporate restrictions imposed in the previous section.

4.1 Data

We use the information provided by the United States Treasury, and the Federal Reserve of Saint Luis Economic Data (FRED). On a daily basis, they publish yields to maturities for Government securities with $1/12$, $1/4$, $1/2$, 1, 2, 3, 5, 7, 10, 20 and 30 years to maturity. These are commonly referred to as Constant Maturity Treasury rates (CMT)[5], and are a result of an interpolation on the closing market bid yields of actively traded securities. We define an estimation window with the same features specified in the previous section. Using the time series of the CMT we bootstrap every observation. Using cubic splines, we interpolate the same terms defined in the other approach, then we construct the spot yield surface. As a result, we have a rates base completely comparable with the one generated in the previous section.

4.2 Estimation

Diebold and Li (2006) noted that λ_t determines the maturity where the factor loading reaches its highest value. They fixed the decay factor at 0.0609, placing the maximum loading at 30 months, which is consistent with the empirical evidence about the curvature. The main advantage of this analysis is that it does not depend on the base you work on. This makes the model easily adaptable to other economies and windows.

[5]These securities assume a semiannual payment. In case of maturities below a year, the Treasury uses recently issued bills and estimates their bond equivalent yield.

Gilli et al (2010) find that for a certain range of values for λ_t the factor loadings of slope and curvature become highly correlated. This problem could translate in collinearity within the model, allowing parameters to exchange with each other at their estimation. If we fit a cross-section model, collinearity is not a problem itself[6]. It becomes an issue when using the model with forecasting purpose, because of the need for stable and identifiable time series of factors.

$$\min_{\lambda} R = \rho^2_{F_S,F_C} \tag{4}$$

$$F_S = \begin{bmatrix} \left(1-e^{-\lambda\tau_1}\right)\left(\lambda\tau_1\right)^{-1} \\ \left(1-e^{-\lambda\tau_2}\right)\left(\lambda\tau_2\right)^{-1} \\ \vdots \\ \left(1-e^{-\lambda\tau_n}\right)\left(\lambda\tau_n\right)^{-1} \end{bmatrix}, \; F_C = \begin{bmatrix} \left(1-e^{-\lambda\tau_1}\right)\left(\lambda\tau_1\right)^{-1} - e^{-\lambda\tau_1} \\ \left(1-e^{-\lambda\tau_2}\right)\left(\lambda\tau_2\right)^{-1} - e^{-\lambda\tau_2} \\ \vdots \\ \left(1-e^{-\lambda\tau_n}\right)\left(\lambda\tau_n\right)^{-1} - e^{-\lambda\tau_n} \end{bmatrix}$$

We intend to fix λ_t using a similiar logic described in the previous paragraph. Instead of focusing on the curvature maximization, we try to work on the correlation problem, exposed by Gilli et al. (2010). In the excercise (4) we look for the λ that minimizes the squared correlation between the slope (F_S) and curvature (F_C) factor loadings. To evaluate the correlation, we create a vector τ containing daily terms to maturity between 30 days up to 30 years. After running the optimization, the decay factor that nullifies the correlation is 0.2262. In order to linearize the estimation of the model, and face its collinearity problem, we fix λ_t at this value[7].

$$\min_{\theta} G = \frac{1}{2}\left\|y_t - \hat{y}_t\right\|_2^2 \tag{5}$$

$$\theta = \{l_t, s_t, c_t\}$$

$$s.t. \begin{cases} 0 < l_t \leq 15 \\ -15 \leq s_t \leq 15 \\ -15 \leq c_t \leq 15 \\ 0 < l_t + s_t \end{cases}$$

In order to measure the performance of the chosen value, we run a historical simulation, over the same window defined in the previous section. We fix λ_t

at 0.2262 and solve the problem (5) by the restricted least squares (RLS) method. Working over the FRED base, constructed previously, the results show

[6] Actually, fitting the model without fixing the decay factor will increase its degrees of freedom.

[7] It is not necessary for the practitioner to repeat this process. This value can be used to linearize the model for any 30-year yield curve, of any economy.

Figure 5: Time series of factors

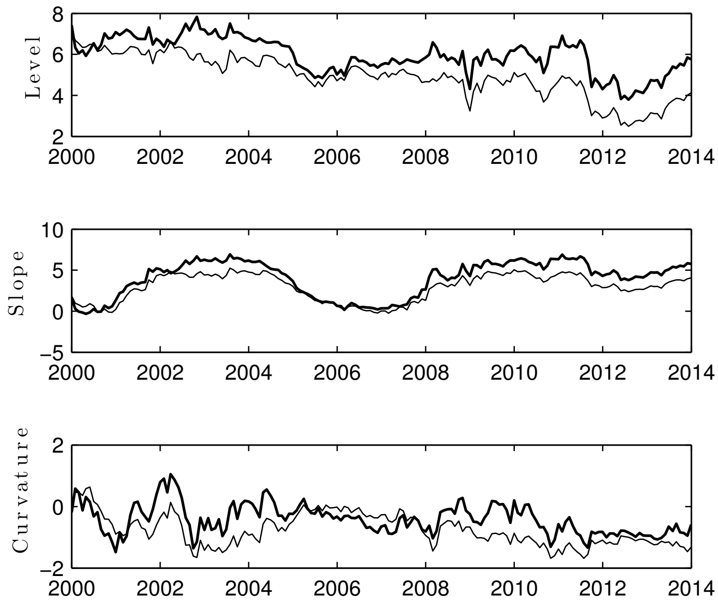

The time series of the factors of level, slope, and curvature are graphed with a thick line. The practial representations of the factors are graphed with a thin line.

that the mean absolute error (MAE) over the whole sample is 8.92 basis points. Figure 5 shows that the estimated time series are highly correlated with the empirical factors, over the estimation window. The level factor has a correlation of 0.85 with the 30-year spot rate. The slope factor shows a correlation of 0.97 with the negative of the slope (measured with the difference between the 3-month and the 30-year spot rates). The curvature factor has a correlation of 0.55 with a butterfly constructed with the 3-month, 2-year and 30-year spot rates. Using the value proposed by Diebold and Li (2006), we obtain a MAE of 12.47 basis points. Our results are satisfactory because by using the decay factor proposed in this section, we are able to achieve a better fit, in contrast to using the value suggested by the original study. Additionally we are able to face the collinearity problem.

5 Conclusions

Modelling the yield curve has never been a trivial process. This is the principal reason it is still one of the most prominent focuses of fixed income research over the decades. Few models have shown the success and popularity of the Nelson-Siegel model. This could be attributed to its parsimonious and intuitive interpretations of its factors (level, slope, and curvature).

The extension that transforms it in a dynamic model has allowed other uses to be discovered. The most known application is the ability to forecast the factors, and hence, the entire yield curve.

This note has described two methodologies to estimate the time series of the factors. The first uses a very rich Treasury securities price base and fits the model using a hybrid optimization method with a genetic algorithm to face the optimization problem. The second methodology presented has a more practitioner approach. It works on a publicly available base, and linearizes the problem fixing the decay factor at the value that minimizes the correlation between the slope and curvature factor loadings. Over an estimation window, with daily frequency, between December of 1999 and December of 2013, both approaches provide a good fit, in terms of mean absolute error (MAE). The rigourous approach gives an average deviation of 8.97 basis points, from market rates. The practitioner approach deviates 8.92 basis points, on average. As a reference point, we estimated the deviaton for both methodologies, fixing λ at the value proposed by Diebold and Li (2006). Our results are satisfactory because both methodologies proposed to fix the decay factor show better fit, in comparison with the original study.

References

[1] Alfaro, R., Becerra, S., and Sagner, A. Estimación de la estructura de tasas nominales en Chile: Aplicación del modelo dinámico de Nelson-Siegel. *Economía Chilena*, 14(3):57–74, 2011.

[2] Christensen, J. H. E., Diebold, F. X., and Rudebusch, G. D. The affine arbitrage-free class of Nelson-Siegel term structure models. *Journal of Econometrics*, (164):4–20, 2011.

[3] Diebold, F. X., Ji, L., and Li, C. A three-factor yield curve model: Non-affine structure, systematic risk sources and generalized duration. *Macroeconomics, Finance and Econometrics: Essay in Memory of Albert Ando*, pages 240–274, 2006.

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

[5] Diebold, F. X. and Rudebusch, G. D. *Yield Curve Modeling and Forecasting: The Dynamic Nelson-Siegel Approach*. Princeton University Press, 2013.

[6] Gilli, M., Grosse, S., and Schumann, E. Calibrating the Nelson-Siegel-Svensson model. *COMISEF Working Paper Series No.31*, 2010.

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

[8] Jones, F. J. Yield curve strategies. *The Journal of Fixed Income*, 1(2):43–48, 1991.

[9] Litterman, R. and Scheinkman, J. Common factors affecting bond returns. *The Journal of Fixed Income*, 1(1):54–61, 1991.

[10] Nelson, C. R. and Siegel, A. F. Parsimonious modeling of yield curves. *Journal of Business*, 60(4):473–489, 1987.

[11] Svensson, L. Estimating forward interest rates with the extended Nelson-Siegel method. *Sveriges Riskbank Quarterly Review*, 3:13–26, 1995.

A Matlab Code

```
%% Parameters.
clear, clc
series = {'DGS3MO' 'DGS6MO' 'DGS1' 'DGS2' 'DGS3' 'DGS5' 'DGS7'...
        'DGS10' 'DGS20' 'DGS30'};
cmt_terms = [3/12 6/12 1 2 3 5 7 10 20 30];
spot_terms = [3/12 6/12 9/12 12/12 15/12 18/12 21/12 24/12 30/12...
        36/12 4:1:30];
window ={'12/31/1999' '12/31/2013'};
tol = 3; % Tolerance for missing data from FRED.
lambda = 0.2262;

%% CMT yields base construction.
% Conection with FRED.
conection = fred('https://research.stlouisfed.org/fred2/');
nseries = length(series); % We count the number of series.
for i = 1:nseries
        get = fetch(conection, series(i), window(1), window(2));
        if i == 1
                cmt_dates = get.Data(:,1);
                cmt_ytm = zeros(length(get.Data),nseries);
        end
        cmt_ytm(:,i) = get.Data(:,2);
end
close(conection)
clear i get nseries conection series window

% We clean the base of missing data observations.
index = sum(isnan(cmt_ytm),2) > tol;
count = 0;
for i = 1 : length(cmt_dates)
        if index(i) == 1
                cmt_ytm(i - count,:) = [];
                cmt_dates(i - count,:) = [];
                count = count + 1;
        end
end
clear i count index tol

%% Bootstrapping the yield surface.
spot_surf = zeros(length(cmt_dates), length(spot_terms));
for i = 1:length(cmt_dates)
        aux = [cmt_ytm(i,:)'/100 zeros(length(cmt_terms),1)];
        for j = 1:length(cmt_terms)
                aux(j,2) = addtodate(cmt_dates(i,1),...
                        cmt_terms(j)*12,'month');
        end
        aux = aux(~any(isnan(aux),2),:);
        [z,t] = pyld2zero(aux(:,1),aux(:,2),cmt_dates(i));
        t = yearfrac(cmt_dates(i),t);
        spot_surf(i,:) = spline(t,z,spot_terms)*100;
end
clc, clear i j z t aux

%% Curve Fitting.
loadings = zeros(length(spot_terms),3);
```

```
loadings(:,1) = 1;
loadings(:,2) = (1-exp(-lambda.*spot_terms'))./ ...
        (lambda.*spot_terms');
loadings(:,3) = (1-exp(-lambda.*spot_terms'))./...
        (lambda.*spot_terms')-exp(-lambda.*spot_terms');
lsc = zeros(length(cmt_dates),3); % Time series vector.

% Optimization.
bounds = [0 -15 -15; 15 15 15]; % Parameters boundries.
rest = [-1,-1,0]; % (b1 + b2) > 0
opt = optimoptions(@lsqlin,'Display','off','Algorithm','active-set');
for i = 1:length(cmt_dates)
        lsc(i,:) = lsqlin(loadings,spot_surf(i,:),rest,0,[],[],...
                bounds(1,:),bounds(2,:),[],opt);
end
```