ERASMUS UNIVERSITY ROTTERDAM
Erasmus School of Economics

Pricing the term structure of government and corporate bonds with linear regressions[1]

Bachelor Thesis Econometrics and Operations Research

Thomas Geissler (445907)
445907tg@student.eur.nl

Supervisor:
Prof. Dr. M. van der Wel

Second assessor:
X. Gong MSc

Date final version:
July 6, 2019

Abstract

In this paper, I study a three-step linear regression model to forecast U.S. Treasury bond yields as well as corporate bond yields. Previous studies applied this model to government bonds which have virtually no risk of default, and found that they could predict yields very accurately. However, no study has applied the model to corporate bonds. I find that five pricing factors obtained using principal component analysis, provide the best forecast for Treasury and corporate bonds. The predicted yields from the model show significant bias for Treasury bonds with higher time-to-maturities, but after correcting this bias, the predicted yields have an almost perfect fit to the actual yields. For the corporate bonds, the model performs reasonably in periods where the economy is stable, but cannot be used for forecasting. In periods with high credit risk, the model does not predict or follow the yields correctly since the model predicts that yields go down instead of up. For larger time-to-maturities the model is accurate when reflecting the yields about the x-axis.

[1]The views stated in this thesis are those of the author and not necessarily those of Erasmus School of Economics or Erasmus University Rotterdam.

Contents

1 Introduction

Traditionally, bond prices are calculated by discounting the cash flows of the bond. The correct discount rate is not a simple rate and has been studied widely and intensively. Inflation, liquidity, and risk of default are some critical aspects for the calculation this discount rate. If an investor has the choice between a bond that matures in one year and a bond that matures in five years, he may want to receive a higher yield for the bond with a time-to-maturity of five years compared to the bond with a shorter maturity. The investor could expect that the company defaults on the loan after two years, and wants to be compensated for taking on that risk. We could also think of these yields as a group, and draw a line through all the yields at that point in time. This line is the yield curve or the term structure of yields. Usually, the curve is upward sloping, but it can also be inverted, or have some bumps. Haubrich and Dombrosky (1996) found that an inverted yield curve could be a sign of an oncoming recession. This is one of the many reasons why forecasting the term structure is essential for investors and governments.

Bonds are assets, like stocks, whose returns do not follow a normal distribution (De Pooter, Ravazzolo, & Van Dijk, 2010; Sheikh & Qiao, 2010). What sets bonds apart from stocks is not only the obligation of the issuer to pay off the loan but also the fact that bonds are traded with different maturities. Changes in the short-term yields lead to changes in the long-term yields. These cross-equation restrictions make it hard to forecast the yields using standard vector autoregressive models. The manipulation of short-term interest rates by the Federal Reserve makes forecasting even more complicated. Affine term structure models can capture these properties. These models require three assumptions: the errors of the state variables are conditionally normal, the pricing kernel is exponentially affine, and the prices of risk are affine. Maximum likelihood is primarily used to estimate the coefficients of the model, see Pearson and Sun (1994) and Ang and Piazzesi (2003). However, maximum likelihood estimation is computationally very expensive and could lead to bias in small samples. Improvements have been made by Hamilton and Wu (2012) to replace maximum likelihood, and by Bauer, Rudebusch, and Wu (2012) to correct for the bias. Joslin, Singleton, and Zhu (2011) show that a vector autoregressive model (VAR) with one lagged term instead of maximum likelihood leads to similar if not identical parameter estimations. However, in the next step, where the price evolution parameters are used to obtain parameters of risk factors, they go back to numerical optimization.

Adrian, Crump, and Moench (2013) introduce a different model, which does not use maximum likelihood at all. Instead, the model uses linear regressions, which is much faster than maximum likelihood. The model can also be used for coupon-bearing bonds, which is another advantage of the use of linear regressions. Furthermore, pricing factors do not have to be a linear combination of yields, but can also be observable factors such as inflation or unemployment rate. The authors show that their model, combined with five pricing factors, fits exceptionally well to the observed yields. The model starts by quantifying the relation between the pricing factors and the lagged pricing factors. Next, the model estimates the influence of shocks and lagged pricing factors on excess returns. The last step involves estimating the price of risk parameters from the previous regression. Using the estimated parameters of these linear regressions, Adrian et al. (2013) reconstruct the yield curve with recursive linear restrictions.

In this paper, I replicate the findings of Adrian et al. (2013) by finding the correct factor specification and by applying the three linear regressions to estimate the yield curve of U.S. Treasury yields. Furthermore, I extend their research by applying the model to credit risk of companies. First, I create corporate bond portfolios, and second, I apply the model introduced by Adrian et al. (2013). It turns out that the estimated yields do not fit the actual yields accurately. Whereas the corporate bond yields go up in the financial crisis, the predicted yields go down. On the other hand, the U.S. Treasury bond yields fit the model very accurately.

Corporate bonds are different from Treasury bonds, for plenty of reasons. The most important reason is that corporate bonds have a risk of default. The issuing company could go bankrupt, and may not be able to pay back the full amount of money owed. The U.S. Treasury has virtually no risk of default since it can print more money to pay off the loan. However, the Treasury has a debt ceiling, so it could potentially run out of money. This has never happened because the United States Congress raises the debt limit when deemed necessary. Because of the risk of default, investors want a higher yield than the government bonds. Rating agencies such as Standard & Poor's assess the creditworthiness of the issuing company. Usually, the lower the credit rating, the higher the bond yield. Assuming that investors want a higher yield from bonds, while not adding too much risk, they might be interested in adding corporate bonds to their portfolio.

The term structure models can also be extended to find risk premia, like the inflation risk premium and the credit risk premium. I apply the model of Adrian et al. (2013) to a portfolio of corporate bonds with different time-to-maturities, and estimate the yield curve for the portfolios. By creating portfolios, I minimize the idiosyncratic risk of the corporate bonds, which results in a more accurate panel of data. Chen, Liu, and Cheng (2010) are able to estimate the inflation and inflation risk premium using a quadratic term structure model but is reliant on maximum likelihood. Follow-up research by Abrahams, Adrian, Crump, Moench, and Yu (2016) extended the model of Adrian et al. (2013) by allowing it to decompose the real and nominal yields using linear regressions. From the difference between the two, they are able to generate an estimate for the inflation risk premium. On the other hand, credit risk premia are much more difficult to obtain. It involves regime switching models, assumptions about credit ratings, assumptions about how much debt can be repaid, and it can get much more complicated by including 'jumping to default' (see Jarrow, Lando, and Turnbull (1996) and Zhou (2001)). However, this is far outside the scope of this thesis.

The rest of this paper is organized as follows: I will start by presenting the three-step model, and how to reconstruct the yields from the resulting parameters in section 2. Next, I describe the data I used, and how I transformed it to use it for the research in section 3. In section 4, I present the results of my research and the interpretation of these results. Finally, section 5 concludes by giving a summary and a discussion.

2 Methodology

The model that Adrian et al. (2013) introduce, which I will refer to as the ACM model, requires a specification of pricing factors. In this section, I present how I obtain these pricing factors, and I give the framework of the ACM model where the factors are used. Next, I show the tests I use to choose the correct number of pricing factors, and finally, I show how I reconstruct the yields using the parameters obtained from the ACM model.

2.1 Pricing factors

In order to obtain pricing factors from the data, I use principal component analysis (PCA). I construct the correlation matrix of a cross-section of yields of $N = 40$ time-to-maturities for the government bonds, and $N = 39$ for the corporate bonds. Next, I apply PCA to this correlation matrix, which results in the eigenvector matrix $\mathbf{E}$. Given $\mathbf{E}$ and the yields, I construct the pricing factors by constructing the principal components,

$$\mathbf{P}_t = \mathbf{E}'\mathbf{R}_t, \tag{2.1}$$

where $\mathbf{R}$ is the matrix with yields. This equation boils down to

$$P_{it} = e_{i1}R_{1t} + e_{i2}R_{2t} + \cdots + e_{iN}R_{Nt}, \tag{2.2}$$

where e_{1i} is the first value of the i-th eigenvector. $\mathbf{P}_t$ is an $N \times 1$ vector, and by choosing the first K components, I have created the pricing vectors for every $t = 1, \ldots, T$.

Litterman and Scheinkman (1991) showed that one could capture enough of the variation over time and over multiple maturities with three factors. These factors are level, slope and curvature, which corresponds to the parameters of the Nelson-Siegel yield curve (Nelson & Siegel, 1987). Later research of Cochrane and Piazzesi (2005) and Joslin, Priebsch, and Singleton (2014) showed that additional factors, such as the "bond-return forecasting factor" and inflation are also crucial. Cochrane and Piazzesi (2009) obtain pricing factors by selecting the first three principal components and then add the return forecasting factor. This factor should predict the one-month excess return better than the fourth and the fifth principal component. Adrian et al. (2013) calculate this forecasting factor by regressing the excess returns on a vector of ten one month lagged one-year forward rates (F_t),

$$rx_{t+1} = \gamma_0 + \Gamma F_t + \eta_{t+1}. \tag{2.3}$$

After this, they apply PCA to $\hat{\Gamma} F_t$ and take the first principal component as the return forecasting factor. The forecasting factor is then added to the three principal components of the original yields, such that

$$P_t = [x_t \; PC1_t \; PC2_t \; PC3_t]', \tag{2.4}$$

where x_t is the forecasting factor, and PC stands for principal component. I will refer to this pricing factor specification as the CP model. Since I do not have forward rates for the corporate bonds, I am not able to estimate the CP model for these types of bonds.

2.2 The model

In order to price the term structure, I use the framework presented by Adrian et al. (2013), and the model described in this section follows their work closely. I present the three steps of the ACM model and the assumptions it is based on. In this section, I assume that the correct number of pricing factors is K.

2.2.1 The first linear regression

I start by estimating a standard vector autoregressive (VAR(1)) model to capture the pricing factor evolution,

$$P_{t+1} = \mu + \Phi P_t + \varepsilon_{t+1}, \tag{2.5}$$

where μ is the $K \times 1$ level parameter, P_t is the $K \times 1$ vector of the pricing factors, Φ is the $K \times K$ coefficient matrix, and ε_t the $K \times 1$ shock or innovation. I assume that the shocks are conditionally normally distributed,

$$\varepsilon_{t+1} | \mathcal{I}_t \sim \mathrm{N}(0, \Sigma), \tag{2.6}$$

where $\mathcal{I}_t$ denotes all historical information of P_t before time t. After performing the regression, I stack the innovations $\hat{\varepsilon}'_{t+1}$ into the matrix $\hat{V}$ and construct the (co)variance matrix $\hat{\Sigma} = \hat{V}'\hat{V}/T$, where T is the number of observations.

2.2.2 The second linear regression

The next regression involves regressing excess returns on a constant, lagged pricing factors, and the pricing factor innovations. I assume that no arbitrage is possible and that a pricing kernel M_t exists, which implies

$$Y_t^{(\tau)} = \mathrm{E}_t \left[M_{t+1} Y_{t+1}^{(\tau-1)} \right]. \tag{2.7}$$

I also assume that this pricing kernel is exponentially affine, such that

$$M_{t+1} = \exp\left(-r_t - \frac{1}{2}\lambda'_t \lambda_t - \lambda'_t \Sigma^{-1/2} \varepsilon_{t+1} \right). \tag{2.8}$$

In this equation, $r_t = \ln Y_t^{(1)}$, which corresponds to the price of the continuously compounded risk-free rate. I further assume that market prices of risk are essentially affine,

$$\lambda_t = \Sigma^{-1/2} \left(\lambda_0 + \lambda_1 P_t \right), \tag{2.9}$$

following Duffee (2002). The excess returns are calculated as

$$rx_{t+1}^{(\tau-1)} = \ln Y_{t+1}^{(\tau-1)} - \ln Y_t^{(\tau)} - r_t, \tag{2.10}$$

where $rx_t^{(\tau-1)}$ is the log excess holding return during time t of a bond maturing in τ months, $\ln Y_t^{(\tau)}$ is the zero-coupon log bond price at time t with time-to-maturity τ. The zero-coupon log bond price can be easily computed with

$$\ln Y_t^{(\tau)} = -\frac{\tau}{12} y_t^{(\tau)}. \tag{2.11}$$

I use the same method of calculating excess returns for the Treasury bonds and corporate bonds, because the yield-to-maturity has the capital gains of the coupons from the corporate bonds already calculated into it. From the yield-to-maturity, I can create a zero-coupon bond, which has a lower price but does not pay out the coupon.

Next, I stack the excess holding returns for all time-to-maturities and over time, and create the $M \times (T-1)$ matrix rx, where M is the number of maturities for which I calculate the excess returns. For the corporate bonds, I take the same risk-free rate as for the Treasury yields, because I assume that an investor can invest in risk-free Treasury bonds.

Furthermore, I assume that the excess returns rx and innovations of equation (2.5) are jointly normally distributed and that the return errors e_{t+1} are independently and identically distributed (iid) with variance σ^2. I use these assumptions to derive the return generating process:

$$rx_{t+1}^{(\tau-1)} = \beta^{(\tau-1)\prime}(\lambda_0 + \lambda_1 P_t) - \frac{1}{2}(\beta^{(\tau-1)\prime}\Sigma\beta^{(\tau-1)} + \sigma^2) + \beta^{(\tau-1)\prime}\varepsilon_{t+1} + e_{t+1}^{(\tau-1)}. \tag{2.12}$$

Here, $\beta^{(\tau-1)}$ can be seen as a factor loading, and risk prices are λ_0 and λ_1. I exclude the full derivation since it is not the focus of this paper; the derivation can be found in Adrian et al. (2013). Equation (2.12) can be stacked for the maturities and over time as

$$rx = \boldsymbol{\beta}' \left(\lambda_0 \iota_T' + \lambda_1 P_-\right) - \frac{1}{2}\left(B^\star \operatorname{vec}(\Sigma) + \sigma^2 \iota_N\right)\iota_T' + \boldsymbol{\beta}'\hat{V} + E, \tag{2.13}$$

and summarised as

$$rx = \mathbf{a}\iota_{T-1}' + \boldsymbol{\beta}'\hat{V} + \mathbf{c}P_- + E, \tag{2.14}$$

where ι_{T-1} is a $(T-1)\times 1$ vector of ones, $P_- = [P_0 \; P_1 \; \dots \; P_{T-2}]$ is a $K \times (T-1)$ matrix of lagged pricing factors, and E an $M \times (T-1)$ matrix of residuals. I collect the regressors in the $(2K+1)\times(T-1)$ matrix

$$\tilde{Z} = \left[\iota_{T-1} \; \hat{V}' \; P_-'\right]', \tag{2.15}$$

and the coefficients are collected in the $M \times (2K+1)$ matrix

$$[\hat{\mathbf{a}} \; \hat{\boldsymbol{\beta}}' \; \hat{\mathbf{c}}] = rx\tilde{Z}'\left(\tilde{Z}\tilde{Z}'\right)^{-1}. \tag{2.16}$$

I can now estimate σ^2 following

$$\hat{\sigma}^2 = \frac{\operatorname{tr}\left(\hat{E}\hat{E}'\right)}{N(T-1)}, \tag{2.17}$$

and $B^\star$ as

$$\hat{B}^\star = [\operatorname{vec}(\hat{\boldsymbol{\beta}}^{(1)}\hat{\boldsymbol{\beta}}^{(1)\prime}) \; \dots \; \operatorname{vec}(\hat{\boldsymbol{\beta}}^{(M)}\hat{\boldsymbol{\beta}}^{(M)\prime})]', \tag{2.18}$$

where $\hat{\boldsymbol{\beta}}^{(i)}$ is the i-th column of $\hat{\boldsymbol{\beta}}$.

2.2.3 The third linear regression

The final regression involves estimating the risk parameters λ_0 and λ_1. From equation (2.13) and equation (2.14), I derive that $\mathbf{a} = \boldsymbol{\beta}'\lambda_0 - \frac{1}{2}\left(B^\star \operatorname{vec}(\Sigma) + \sigma^2\iota_N\right)$ and $\mathbf{c} = \boldsymbol{\beta}'\lambda_1$. I use these expressions to derive the following estimators of λ_0 and λ_1 using a cross-sectional regression:

$$\widehat{\lambda}_0 = (\hat{\boldsymbol{\beta}}\hat{\boldsymbol{\beta}}')^{-1}\hat{\boldsymbol{\beta}}\left(\hat{\mathbf{a}} + \frac{1}{2}\left(\hat{B}^{\star}\operatorname{vec}(\hat{\Sigma}) + \hat{\sigma}^2\iota_M\right)\right), \tag{2.19}$$

and

$$\widehat{\lambda}_1 = \left(\hat{\boldsymbol{\beta}}\hat{\boldsymbol{\beta}}'\right)^{-1}\hat{\boldsymbol{\beta}}\hat{\mathbf{c}}. \tag{2.20}$$

Adrian et al. (2013) go on to show that the variance of $\boldsymbol{\beta}$ is

$$\mathcal{V}_\beta = \sigma^2\left(I_M \otimes \Sigma^{-1}\right), \tag{2.21}$$

where $\otimes$ is the Kronecker product.

2.3 Model specification test

To analyze the correct number of pricing factors, I perform two model specification tests. Adrian et al. (2013) argue that the matrix of factor loadings $\boldsymbol{\beta}$ should be of full column rank to identify λ_0 and λ_1. Therefore, the first test is the Anderson (1951) canonical correlations test, which tests whether the rank of a matrix is r or that it is not greater than a specified value p. Under the null that $\operatorname{rank}(\boldsymbol{\beta}) \leq r < K$, the Anderson statistic is

$$rk_r = -(T-1)\sum_{i=r+1}^{K}\ln\left(1-\rho_i^2\right) \overset{a}{\sim} \chi^2((K-r)(M-r)), \tag{2.22}$$

where ρ_i is de sample partial canonical correlation of V on rx conditional on X_-. For a K-factor model, I test whether $r = K-1$, such that the equation collapses to

$$rk_r = -(T-1)\ln\left(1-\rho_K^2\right) \overset{a}{\sim} \chi^2((M-K+1)). \tag{2.23}$$

The second test is a test for useless factors by testing whether certain columns of $\boldsymbol{\beta}$ are equal to zero using the Wald test. Under the null that $\beta_i = 0_{M\times 1}$, the statistic is

$$W_{\beta_i} = \hat{\beta}_i'\hat{\mathcal{V}}_{\beta_i}^{-1}\hat{\beta}_i \sim \chi^2(M), \tag{2.24}$$

where β_i is the i-th column of $\boldsymbol{\beta}'$ and $\mathcal{V}_{\beta_i}$ is the variance of β_i. In a K-factor model, I test whether the K-th column of $\boldsymbol{\beta}'$ is zero. The significance level for both tests is set at $\alpha = 0.05$. I perform these tests on four different models: three ACM models with pricing factors equal to three, four or five, and the CP model. For clarity, I should note that I test whether the third principal component in the CP model is zero. As I indicated before, I do not construct the CP model for the corporate bonds, so I only perform the tests on the ACM models.

2.4 Constructing the curve

I am able to generate a zero-coupon yield curve with the parameters $\{\Phi, \Sigma, \boldsymbol{\beta}, \sigma, \lambda_0, \lambda_1\}$ from section 2.2. Using the assumptions, the bond prices can be shown as

$$\ln Y_t^{(\tau)} = A_\tau + B_\tau' P_t + u_t^{(\tau)}. \tag{2.25}$$

Substituting equation (2.25) into equation (2.10) and setting it equal to equation (2.12), I get

$$\begin{aligned} A_{\tau-1} + B_{\tau-1}'(\mu + \Phi P_t + v_{t+1}) + u_{t+1}^{(\tau-1)} - A_\tau - B_\tau' P_t - u_t^{(\tau)} + A_1 + B_1' P_t + u_t^{(1)} \\ = \beta^{(n-1)\prime}(\lambda_0 + \lambda_1 P_t + v_{t+1}) - \frac{1}{2}(\beta^{(\tau-1)\prime}\Sigma\beta^{(\tau-1)} + \sigma^2) + e_{t+1}^{(\tau-1)}. \end{aligned} \tag{2.26}$$

This equation has to hold for all $t = 1, \ldots, T$. If I set $A_1 = -\delta_0$ and $B_1 = -\delta_1$, the following set of recursive restrictions for the parameters A and B has to hold:

$$\begin{cases} A_\tau = A_{\tau-1} + B_{\tau-1}'(\mu - \lambda_0) + \frac{1}{2}(B_{\tau-1}'\Sigma B_{\tau-1} + \sigma^2) - \delta_0, \\ B_\tau' = B_{\tau-1}'(\Phi - \lambda_1) - \delta_1', \\ A_0 = 0, \quad B_0' = 0, \\ \beta^{(\tau)} = B_\tau' \end{cases} \tag{2.27}$$

The parameters δ_0 and δ_1 can be estimated by regressing the price of the one-month bonds on the pricing factors, since the equation is

$$\ln Y_t^1 = A_1 + B_1' P_t + u_t^1. \tag{2.28}$$

After obtaining the full vector A, and matrix B, I reconstruct all log bond prices. Following the reconstruction, I convert the prices back into continuously compounded zero-coupon yields by

$$\hat{y}_t^{(\tau)} = -\frac{\ln \hat{Y}_t^{(\tau)}}{\tau/12}. \tag{2.29}$$

The estimated excess returns are shown as

$$\widehat{rx}_{t+1} = \hat{B}_{\tau-1}'\left(\hat{\lambda}_0 + \hat{\lambda}_1 P_t\right) - \frac{1}{2}\left(\hat{B}_{\tau-1}'\hat{\Sigma}\hat{B}_{\tau-1} + \hat{\sigma}^2\right) + \hat{B}_{\tau-1}'\hat{\varepsilon}_{t+1}, \tag{2.30}$$

which corresponds to the return generating process of equation (2.12), where the time-to-maturities are stacked into the vector rx_{t+1}.

3 Data

In this section, I present the data I use to obtain results from the model in the previous section. I start by showing the government bond data, in particular I show how to construct the cross-section of data using the Nelson-Siegel-Svensson model, introduced by Nelson and Siegel (1987) and extended by Svensson (1994). Next, I present the data for the corporate bonds, and how I cleaned it to make maturity sorted portfolios.

3.1 Treasury yield data

I estimate the K-factor specification for the ACM model using the dataset of daily zero-coupon yields which is based on fitted Nelson-Siegel-Svensson curves. Gürkaynak, Sack, and Wright (2007) have constructed these daily yields for United States Treasury Notes, Bills, and Bonds with time-to-maturities (τ) of `1, 2,..., 30` years. Gürkaynak et al. (2007) also included the parameters of the Nelson-Siegel-Svensson curves which I use to make a cross section of yields for $\tau =$ `3, 6,..., 120` months.

Six parameters influence the Nelson-Siegel-Svensson curves: $\{\beta_0, \beta_1, \beta_2, \beta_3, \tau_1, \tau_2\}$. When these parameters are given, the zero-coupon yield curve at time t is easily calculated using

$$y_t(\tau) = \beta_0 + \beta_1 \frac{1 - \exp\left(-\frac{\tau}{\tau_1}\right)}{\frac{\tau}{\tau_1}} + \beta_2 \left[\frac{1 - \exp\left(-\frac{\tau}{\tau_1}\right)}{\frac{\tau}{\tau_1}} - \exp\left(-\frac{\tau}{\tau_1}\right)\right] + \beta_3 \left[\frac{1 - \exp\left(-\frac{\tau}{\tau_2}\right)}{\frac{\tau}{\tau_2}} - \exp\left(-\frac{\tau}{\tau_2}\right)\right], \tag{3.1}$$

where τ is the time-to-maturity, and each of the Nelson-Siegel-Svensson parameters should have index t, but I omit them since this would make it harder to read. To calculate the return forecasting factor of the CP model, which is discussed in section 2.1, I construct one-year forward rates with the Nelson-Siegel-Svensson parameters. I do this with the following equation:

$$f_t(\tau, 0) = \beta_0 + \beta_1 \exp\left(-\tau/\tau_1\right) + \beta_2 \left(\tau/\tau_1\right) \exp\left(-\tau/\tau_1\right) + \beta_3 \left(\tau/\tau_2\right) \exp\left(-\tau/\tau_2\right). \tag{3.2}$$

In section 2.2.2, I defined the price of the continuously compounded risk-free rate as $r_t = \ln Y_t^{(1)}$, and I calculate this price using equation (2.11). The zero-coupon yield of this risk-free bond is also calculated using the Nelson-Siegel-Svensson model. I set $\tau = 1$ in equation (3.1), such that the risk-free rate is based on the one month yield.

The data on the parameters spans from November 25, 1985, to April 26, 2019. The dataset contains information of the parameters before November 25, 1985, but not all bonds with a longer maturity were available before then. Therefore, Gürkaynak et al. (2007) highly recommend focusing only on securities that were available during that time. Since the model of Adrian et al. (2013) is based on monthly observations, I take end-of-month values for the parameters.

3.2 Corporate bond data

I extend the research of Adrian et al. (2013) by applying their model to corporate bond data. The data I collect is from the WRDS Bond Returns dataset[2]. It contains all United States monthly corporate bond transactions from July 2002, to March 2018. The data is collected by WRDS from the Trade Reporting and Compliance Engine (TRACE) of the Financial Industry Regulatory Authority (FINRA), and the Mergent Fixed Income Securities Database (FISD). The data is essentially pre-cleaned by WRDS since they already removed about 28% of the raw data from TRACE and FISD. The cleaned dataset contains information of the issuer, the credit

[2] Available at `https://wrds-www.wharton.upenn.edu/pages/grid-items/wrds-bond-returns/`

rating, the type of bond, when a coupon was not paid, and various other criteria.

From the WRDS Bond Returns dataset, I select six variables: the date, the company's ticker, the Standard & Poor's credit rating, its numerical credit rating, the yield-to-maturity, and the time-to-maturity. With these variables, I construct a cross-section of portfolios, sorted by maturity and time. Before I start constructing the portfolios, I remove all non-investment grade bonds. That is, all bonds with a credit rating of BB+ and worse are deleted to remove the most substantial chunk of the risky bonds. Next, I remove all bonds with a negative yield-to-maturity, because the portfolios are fictive, and therefore it does not have to be very liquid, which is a reason to hold bonds with a negative yield. Another reason to hold bonds with negative yields is that foreign investors might expect the domestic currency to rise, but this is beyond the scope of this paper; thus, the negative yield bonds are ignored.

After cleaning the data, I create $N = 39$ portfolios with $\tau =$ `6, 9, ..., 120` months at each point in time. To do this, I take all the bonds with time-to-maturity between $(\tau - 6, \tau + 6)$ months, and calculate the average yield-to-maturity, price, and coupon. Figure 1 shows the yield-to-maturity of all portfolios over time. The financial crisis of 2008-2009 is very well captured. We can see that short-term bonds have a very high yield, indicating an increase in credit risk and liquidity risk.

Whereas the Treasury yields are zero-coupon yields, the yields of these portfolios are coupon-bearing bond portfolios. However, this should not be a problem since Adrian et al. (2013) argue that only two inputs for their model are necessary: a set of excess returns and pricing factors.

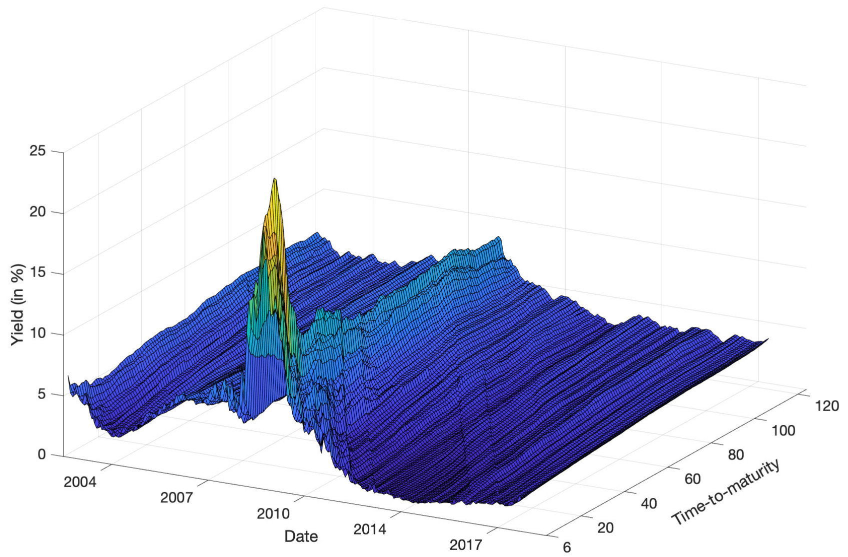

Figure 1: Yields of corporate bonds over all time-to-maturities and time. The yields are calculated by averaging all U.S. corporate bond yields with a maturity within six months of the desired time-to-maturity τ, as described in section 3.2.

4 Results

In this section, I present the results of the ACM model with the data I collected. I start with the results of the Treasury bonds, and finish with the corporate bonds. I show the results of the principal component analysis, the outcome of the model specification tests, and the fit of the chosen factor specification.

4.1 Treasury bonds

For the U.S. Treasury bonds, I use zero-coupon yields as described in section 3.1. I have a cross-section of $N = 40$ time-to-maturities over $T = 402$ months. The first three principal components extracted from this panel, correspond to the level, slope, and curvature factors; the fourth corresponds to the relation between the short and medium maturities, and the fifth to the relation between the short, short-to-medium and long maturities. The ACM model with $K = 5$ factors fits the yields the best, according to the specification tests. The model does have some bias, but after correcting this, the largest yield fitting error is 0.095%.

4.1.1 Principal components and pricing factors

After performing principal component analysis on the correlation matrix of the yields, I selected the first five principal components, which can be seen in figure 2. In figure 2a, the expected three factors of level, slope, and curvature can be found. The first principal component corresponds to the level as it stays relatively constant for all maturities. The second principal component corresponds to the slope, and the third component is the curvature factor. The fourth and fifth principal component can be seen in figure 2b. The fourth component highlights the relation between short and medium-term bonds, and the fifth component can be interpreted as the relation between short, short-to-medium, and long term bonds. The fourth and fifth component only explain 0.014% of the cross-sectional variation.

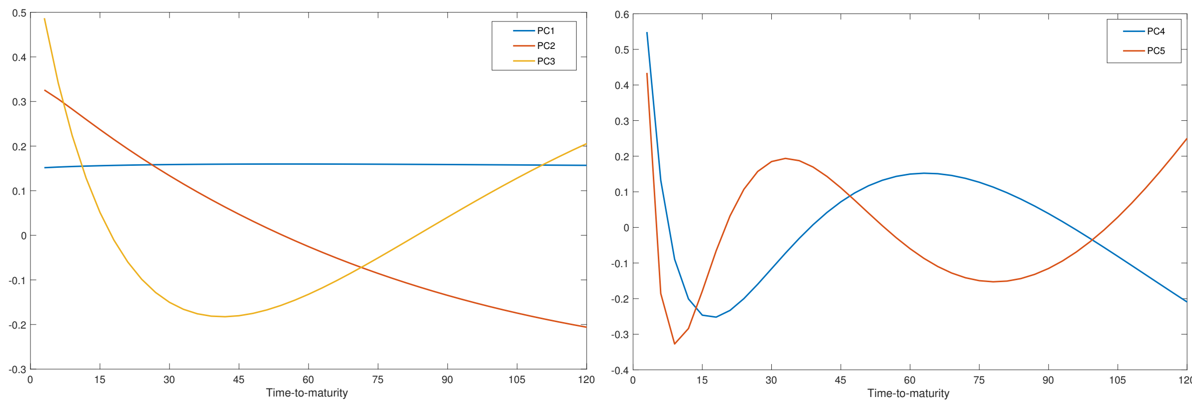

(a) First three principal components **(b)** Last two principal components

Figure 2: Eigenvector loadings for first five principal components of Treasury yields. These plots show the eigenvectors resulting from principal component analysis on the correlation matrix of the panel of zero-coupon yields. The yields are constructed using Nelson-Siegel-Svensson curves, and they span a period of 402 months.

4.1.2 Model specification tests

After obtaining the parameters from the ACM and the CP model, I perform two model specification tests to test the number of pricing factors. The two tests are the Anderson canonical correlations test and the Wald test. Table 1 shows the results of the tests. The ACM model with $K = 5$ pricing factors is selected when I use the Anderson test, while the CP model is selected when I use the Wald test. However, the CP model is the worst in the Anderson test. Therefore, I select the ACM model with $K = 5$ pricing factors as the correct specification.

Table 1: Identification tests for number of pricing factors for Treasury yields.
This table reports the model specification tests that are described in section 2.3. I test the ACM with three, four, and five pricing factors, and the CP model. The first test is the Anderson canonical correlations test, where the statistic rk_{K-1} tests whether the rank of $\boldsymbol{\beta}$ is smaller or equal to $K-1$ in a K factor model. The statistic follows a $\chi^2(M - K + 1)$ distribution, where M is the size of the cross-section of excess returns. The second test is the Wald test, where the statistic W_β tests whether the K-th column of $\boldsymbol{\beta}'$ is equal to $0_{M\times1}$. The statistic follows a $\chi^2(M)$ distribution with significance level at 5%. The p-values of each test is shown in parentheses.

Model	rk_{K-1}	W_β
ACM, $K = 3$	542.06	523.34
	(0.000)	(0.000)
ACM, $K = 4$	562.77	595.35
	(0.000)	(0.000)
ACM, $K = 5$	654.60	1047.35
	(0.000)	(0.000)
CP ($K = 4$)	518.39	3245.74
	(0.000)	(0.000)

From table 1, we can conclude that all models are correctly specified since all p-values are 0.000. We can also see that the ACM model becomes better specified when adding pricing factors. The test statistics consistently increase as K increases, especially from the $K = 4$ to the $K = 5$ model where the Wald statistic nearly doubles. Finally, the CP model performs very well when using the Wald test. This means that the last column of $\boldsymbol{\beta}'$ is significantly different from zero. However, the model performs the worst in the Anderson test. This result means that there is a column in $\boldsymbol{\beta}$ that is spanned by the other columns. Upon further inspection, I find that the return forecasting factor and the second principal component are able to capture a part of the movement of the first principal component. Figure 3 shows this result. Thus, we can conclude that the CP model is not optimally specified.

4.1.3 Fitting the model

Given the five-factor specification for the pricing factors, I construct the forecasts recursively following section 2.4. In figure 4a, the fitted yields with $\tau = 6$ months follow the actual yields very closely, but in figure 4b the fitted yields with $\tau = 120$ months are consistently lower. It turns out that the difference is constant over time, which indicates a presence of bias. This bias might be due to a faulty or incomplete model. Since the bias increases as I increase the time-to-maturity, the fault lies in the recursion of A_τ and B_τ. One of the parameters needed for the recursion is most likely not correct.

I correct this bias by taking the expected yields for the actual and fitted yields and taking

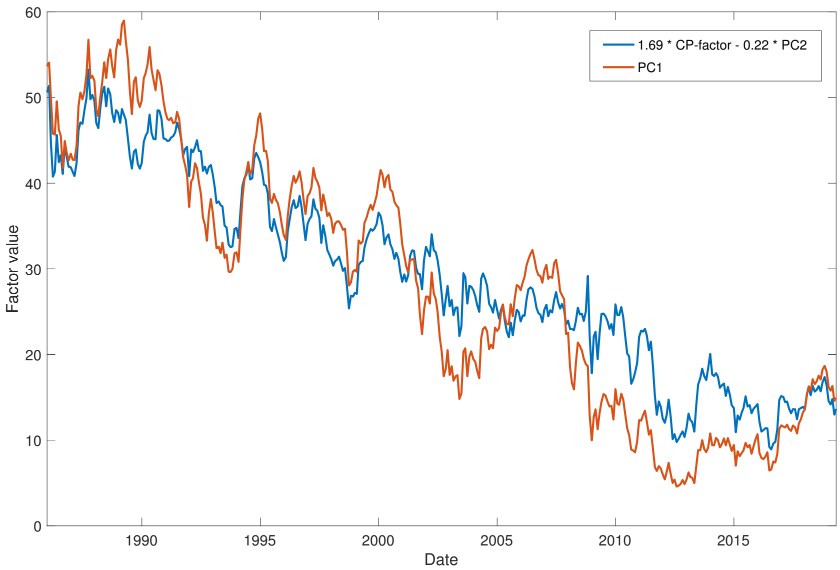

Figure 3: First PC and a linear combination of the return forecasting factor and the second PC. This figure shows that in the CP model, the first principal component of the zero-coupon yields is largely captured by a linear combination of the CP return forecasting factor and the second principal component.

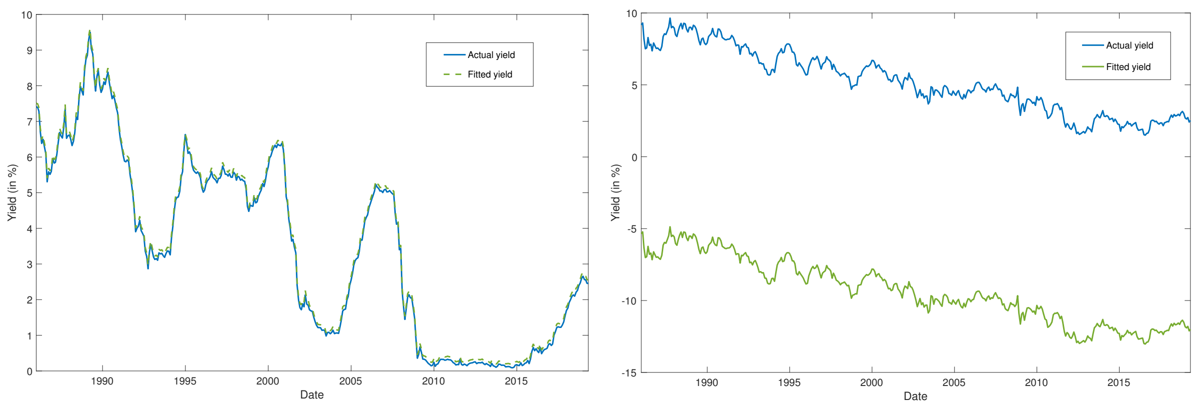

(a) $\tau = 6$ months

(b) $\tau = 120$ months

Figure 4: Fitted yields for the $K = 5$ factor model for different time-to-maturities. This figure provides an overview of the fitted yields obtained from the five factor ACM model. The solid blue line corresponds to the observed yield, the dashed green line in the left figure and the solid green line in the right figure correspond to the fitted yield.

the difference between the two. This can be formalized as

$$b^{(\tau)} = \mathrm{E}\left[y^{(\tau)}\right] - \mathrm{E}\left[\hat{y}^{(\tau)}\right]. \tag{4.1}$$

I do the same bias correction for the excess returns since there is a bias for the higher maturities, as I mentioned before. Figure 10 in appendix A shows the increasing bias of the yields and excess returns as the time-to-maturity increases.

For each τ for which I forecast the yields, I correct the bias following

$$\tilde{y}_t^{(\tau)} = \hat{y}_t^{(\tau)} + b^{(\tau)}, \quad \forall t = 1, \ldots, T. \tag{4.2}$$

After this correction, the fitted yields follow the actual yields much closer as figure 5a and figure 5b present. The biggest error of the yield is only 9.5 basis points or 0.095%. This error occurs in December 1994, for the bonds with a time-to-maturity of 6 months. The largest error of the other time-to-maturities is only two basis points. We can conclude that this model predicts the yields extremely accurately after the bias correction. Figure 5c and figure 5d show the fitted excess returns, and they follow the actual excess returns very closely as well. The largest excess return error is 0.30% across all the time-to-maturities, which highlights the accuracy of the model.

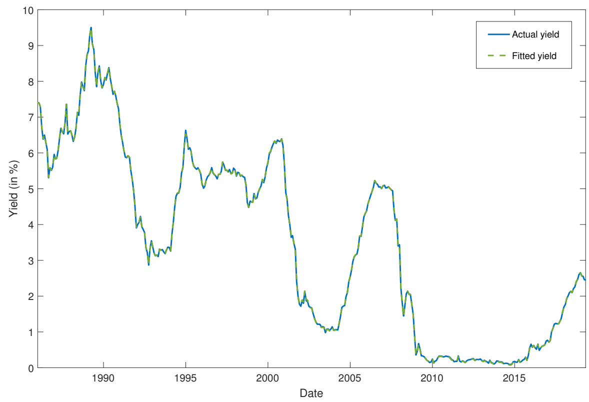

(a) Fitted yields with $\tau = 6$ months

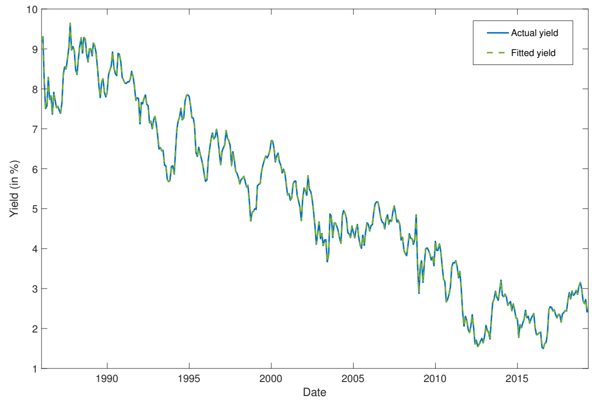

(b) Fitted yields with $\tau = 120$ months

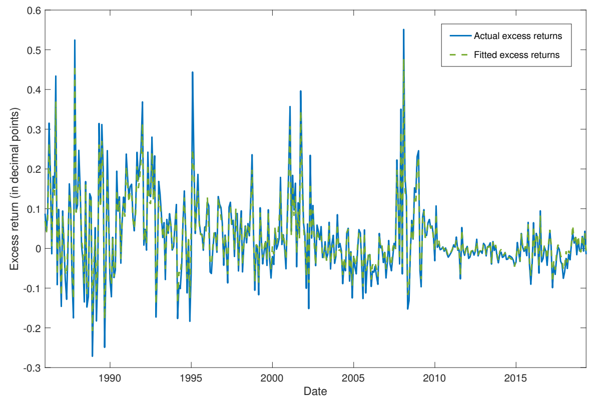

(c) Fitted excess returns with $\tau = 6$ months

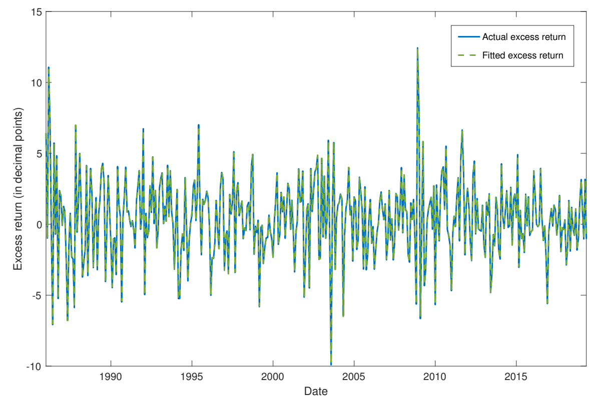

(d) Fitted excess returns with $\tau = 120$ months

Figure 5: Bias-corrected fitted Treasury yields and excess returns for the $K = 5$ factor model for different time-to-maturities. This figure presents plots of fitted yields and excess returns of the five-factor ACM model after a bias correction. The solid blue line is the observed yield and excess return, the dashed green line is the fitted yield and excess return.

Table 2 shows the properties of the forecasting error. The yield errors $\hat{u}$ and excess returns errors $\hat{e}$ are calculated by subtracting the fitted value from the actual value. From this table, we can conclude that the five-factor model fits the actual yields very closely. The standard deviation is tiny; it does not exceed 0.01 for all zero-coupon bonds with a maturity longer than 12 months and is only slightly more than one basis point for $\tau = 6$ months. Finally, consistent with the theory and results of Adrian et al. (2013), I find that the yield errors have a high autocorrelation. We can also conclude that the average return error and standard deviation is very small. In general, the return errors do not show significant autocorrelation, which is in accordance with the results of Adrian et al. (2013).

Table 2: Yield and excess return fitting error properties for different time-to-maturities of Treasury yields.
This table provides a summary of the fitted yield and excess return errors obtained from the five-factor ACM model. The upper panel summarizes the yield errors and the bottom panel the excess return errors. The sample period is December 1985, to April, 2019. The statistics considered are the mean, standard deviation, skewness, kurtosis, and the first and sixth order autocorrelations, $\rho(1)$ and $\rho(6)$ respectively. τ corresponds to the time-to-maturity in months.

Summary statistic	$\tau = 6$	$\tau = 24$	$\tau = 36$	$\tau = 48$	$\tau = 60$	$\tau = 120$
Yield errors ($\hat{u}$)						
Mean	0.000	0.000	0.000	0.000	0.000	0.000
Std. Dev.	0.012	0.005	0.004	0.005	0.004	0.007
Skewness	0.978	0.439	0.066	-0.031	0.203	-0.077
Kurtosis	14.308	3.364	2.686	2.832	2.661	3.104
$\rho(1)$	0.565	0.746	0.815	0.842	0.933	0.835
$\rho(6)$	0.301	0.390	0.695	0.595	0.714	0.499
Excess return errors ($\hat{e}$)						
Mean	0.000	0.000	0.000	0.000	0.000	0.000
Std. Dev.	0.024	0.029	0.029	0.028	0.024	0.045
Skewness	1.228	-0.150	0.211	0.342	0.218	0.214
Kurtosis	14.437	3.565	4.291	5.703	4.819	9.827
$\rho(1)$	0.052	0.124	0.021	-0.094	-0.032	-0.167
$\rho(6)$	0.082	-0.048	-0.043	-0.024	-0.034	-0.022

In general, the kurtosis of the return errors is higher than the kurtosis of the yield errors. This result indicates that the return errors are more heavily tailed. Table 2 shows another interesting result: the kurtosis of the errors with $\tau = 6$ months is significantly higher in comparison to the other time-to-maturities. This result shows that the ACM model produces very small errors, but for the smallest time-to-maturity, large errors tend to be even larger errors compared to other time-to-maturities.

4.2 Corporate bonds

For the corporate bonds, I use the yield-to-maturity of all U.S. corporate bonds, as described in section 3.2. I have a cross-section of $N = 39$ time-to-maturities over $T = 188$ months. The first three principal components are very similar to the first three of the Treasury yields; the fourth and fifth component corresponds to the relation between short and medium-term bonds. The ACM model with $K = 3$ or $K = 5$ factors fits the yields the best, according to the specification tests. The ACM model produces yields with significant bias. Furthermore, the yields tend to move in the opposite direction as the observed yield for longer maturities. For shorter maturities, the predicted yields follow the observed yields only in periods of economic certainty. The model is not able to predict excess returns as the fitted returns are between 8 to 50 times less volatile compared to the observed returns.

4.2.1 Principal components and pricing factors

Figure 6 shows the eigenvectors that resulted from principal component analysis on the correlation matrix of the corporate bond yields. In figure 6a, we can see the three factors of level, slope, and curvature, albeit more roughly than in figure 2a. The first principal component stays relatively constant at 0.15, thus indicating it is a level factor. The second principal component has a downward slope but is much rougher than the second principal component of the zero-coupon yields. The third principal component corresponds to the curvature factor since it goes down at first and turns upwards at $\tau = 30$. The fourth component highlights the relation between short term bonds and short-to-medium term bonds, and the fifth component can be interpreted as the relation between short and medium-to-long term bonds.

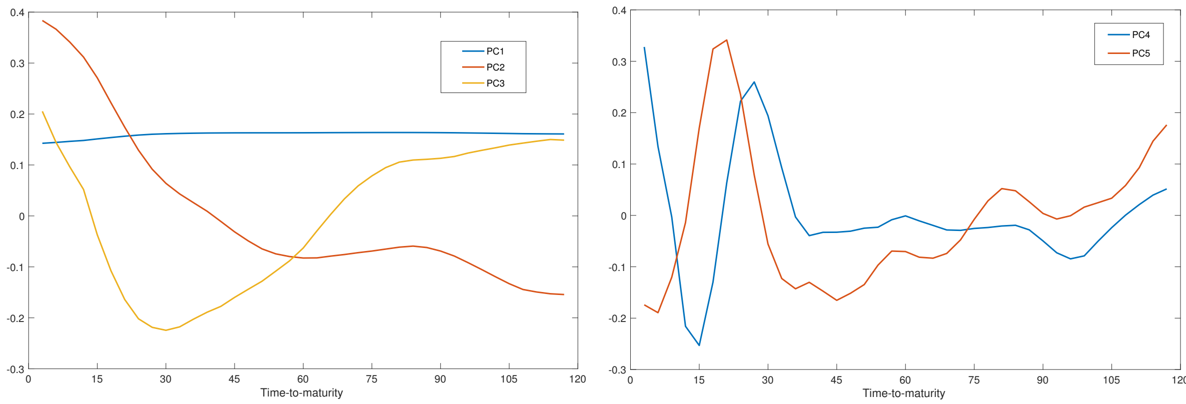

(a) First three principal components

(b) Last two principal components

Figure 6: Eigenvector loadings for first five principal components of corporate bond yields. These figures provide an overview of the eigenvectors obtained from principal component analysis on the correlation matrix of the cross-section of corporate bond yield. The yields are constructed following section 3.2 and span 188 months.

4.2.2 Model specification tests

Table 3 shows the results of the model specification tests for the corporate bonds. The results do not point out that one model works the best. The $K = 5$ model performs the worst in the Anderson test, but the best in the Wald test for useless factors. On the other hand, the $K = 3$ model performs the best in the Anderson test, but second in the Wald test. The $K = 4$ model is the worst in the Wald test and only slightly better than the $K = 5$ model in the Anderson test. Therefore, I select both the $K = 3$ and $K = 5$ factor model as the correct specification for the ACM model, and I keep comparing the two for advantages of one over the other.

Table 3 indicates that only the $K = 4$ factor model is incorrectly specified in the Wald test. However, this model is correctly specified when looking at the Anderson canonical correlation

Table 3: Identification tests for number of pricing factors for corporate bonds.
This table reports the model specification tests that are described in section 2.3. I test the ACM with three, four, and five pricing factors. The first test is the Anderson canonical correlations test, where the statistic rk_{K-1} tests whether the rank of $\boldsymbol{\beta}$ is smaller or equal to $K-1$ in a K factor model. The statistic follows a $\chi^2(M-K+1)$ distribution, where M is the size of the cross-section of excess returns. The second test is the Wald test, where the statistic W_β tests whether the K-th column of $\boldsymbol{\beta}'$ is equal to $0_{M\times 1}$. The statistic follows a $\chi^2(M)$ distribution. The p-values of each test is shown in parentheses.

Model	rk_{K-1}	W_β
ACM, $K=3$	205.467	23.793
	(0.000)	(0.022)
ACM, $K=4$	139.888	12.782
	(0.000)	(0.385)
ACM, $K=5$	134.191	26.706
	(0.000)	(0.009)

test. This means that, although $\boldsymbol{\beta}'$ is full rank, the last column is not significantly different from zero for the four-factor model. Furthermore, the Anderson test shows that the model is better specified when removing pricing factors, which is in contrast with the tests of the Treasury yields. This increase of the statistic when removing pricing factors is especially evident when going from four to three factors. This provides further evidence that the four-factor model performs the worst. In general, the test statistics of the corporate bond yields are lower than the statistics of the Treasury yields. We can conclude that the ACM model works better with Treasury yields than with corporate bond yields.

4.2.3 Fitting the model

It turns out that the ACM model with $K=5$ factors cannot accurately predict the yields for corporate bonds. The yields show a large amount of bias, as figure 7 shows. Figure 11 in appendix B shows the predicted yields of the three-factor model. This model also has a significant bias.

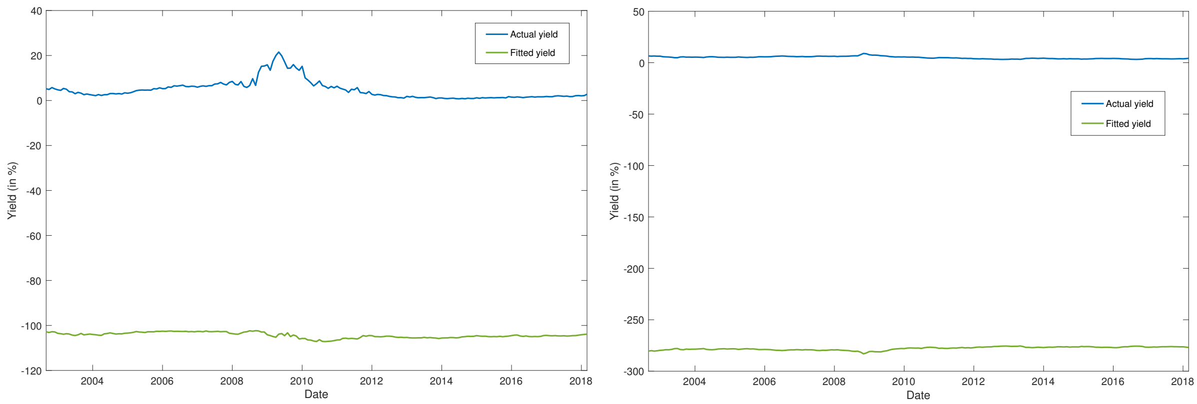

(a) $\tau = 6$ (b) $\tau = 120$

Figure 7: Fitted corporate bond yields for the $K=5$ factor model for different time-to-maturities. The plots in this figure show the fitted yields of corporate bonds obtained from the five factor ACM model. The solid blue line is the observed yield, the solid green line is the fitted yield.

I continue by correcting the bias in the same fashion as in section 4.1.3. Figure 8 shows the results of the bias correction for the five-factor model, and figure 12 in appendix B shows the results for the three-factor model. From figure 8a, we can conclude that the model still does not accurately predict the corporate bond yields with a short time-to-maturity. Whereas yields go up in the financial crisis around 2009, the predicted yields go down. However, from 2014 onward, the predicted yields follow the actual yields better, although they are at a higher level. Figure 8b shows the predicted yields with a time-to-maturity of 120 months. At first sight, we can conclude that the model does not fit the model at all. In fact, it seems to move in the exact opposite direction. The model with $\tau = 120$ months is also where we can see the difference between the three and five-factor ACM model. The $K = 5$ factor model predicts yields go down in 2009, while the $K = 3$ factor model correctly predicts it goes up. However, outside the period of the financial crisis, the yields still seem to move in the opposite direction. In figure 8c and figure 8d, the fitted excess returns are plotted against the observed excess returns. The fitted values are less volatile than the observed excess returns, and do not have any predictive power. For $\tau = 6$ months, the predicted excess returns are fifty times less volatile, and for $\tau = 120$ months, they are eight times less volatile.

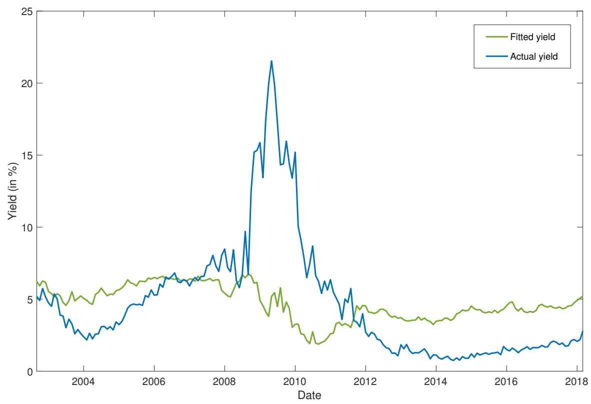

(a) Fitted yields with $\tau = 6$ months

(b) Fitted yields with $\tau = 120$ months

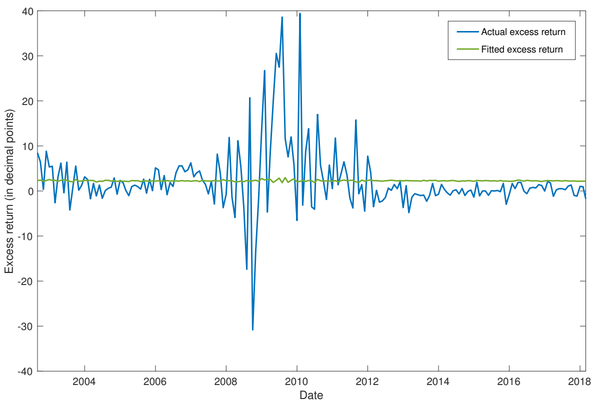

(c) Fitted excess returns with $\tau = 6$ months

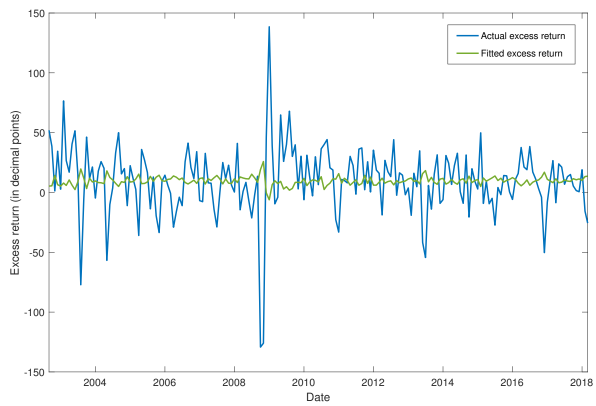

(d) Fitted excess returns with $\tau = 120$ months

Figure 8: Bias corrected fitted corporate bond yields and excess returns for the $K = 5$ factor model for different time-to-maturities. This figure presents plots of fitted yields and excess returns of the five-factor ACM model for corporate bonds after a bias correction. The solid blue line is the observed yield and excess return, the solid green line is the fitted yield and excess return.

Table 4 shows the properties of the forecasting errors. The average yield error is 0.000 by construction, but the standard deviation is extremely high. The errors are also heavily skewed to the right, indicating a lot of positive errors. Most errors are positive because the larger errors in 2008-2010 are negative, and due to the bias correction, the other yields are shifted above the actual yields. The errors show very significant autocorrelation, which is consistent with the results of Adrian et al. (2013) and the results of the Treasury yields. The standard deviation of the excess return errors are relatively small for the bonds with a smaller time-to-maturity compared to the extremely high standard deviation for the longer time-to-maturities. The skewness of the return errors is different from maturity to maturity; in four of the six cases, the skewness is negative. The kurtosis of the errors is extremely high for all maturities, which means that there are large outliers. Finally, in contrast to Adrian et al. (2013), the return errors show significant serial correlation for the first autocorrelation. For $\tau = 6$ and $\tau = 36$ months, the sixth order autocorrelation is also significantly different from zero.

Table 4: Yield and return fitting error properties for different time-to-maturities of corporate bonds. This table provides a summary of the fitted yield and excess return errors obtained from the five-factor ACM model. The upper panel summarizes the yield errors and the bottom panel the excess return errors. The sample period is July 2002, to March, 2018. The statistics considered are the mean, standard deviation, skewness, kurtosis, and the first and sixth order autocorrelations, $\rho(1)$ and $\rho(6)$ respectively. τ corresponds to the time-to-maturity in months.

Summary statistic	$\tau = 6$	$\tau = 24$	$\tau = 36$	$\tau = 48$	$\tau = 60$	$\tau = 120$
Yield errors ($\hat{u}$)						
Mean	0.000	0.000	0.000	0.000	0.000	0.000
Std. Dev.	4.129	2.669	3.054	3.546	3.843	2.747
Skewness	2.045	1.892	2.145	1.971	1.700	0.519
Kurtosis	6.778	6.520	8.618	7.811	6.815	2.687
$\rho(1)$	0.965	0.971	0.970	0.974	0.976	0.972
$\rho(6)$	0.774	0.752	0.707	0.733	0.744	0.812
Excess return errors ($\hat{e}$)						
Mean	0.000	0.000	0.000	0.000	0.000	0.000
Std. Dev.	7.330	11.977	15.690	17.449	18.634	31.900
Skewness	1.724	-0.830	-0.612	0.583	-1.002	-0.802
Kurtosis	12.981	13.577	12.658	13.024	9.731	9.299
$\rho(1)$	0.197	0.162	0.244	0.121	0.300	0.268
$\rho(6)$	0.201	-0.042	-0.119	-0.062	0.013	-0.028

In general, the fitting errors for corporate bonds are higher and more extreme. The return errors show autocorrelation for corporate bonds, but not for Treasury bonds. The errors of $\tau = 120$ months show less variance, skewness, and kurtosis than the other maturities. This is a surprising result since figure 8b shows that the yields seem to move in the opposite direction.

Because the fitted yields in figure 8b move in the opposite direction of the observed yields, I proceed by changing the way I price the bonds in equation (2.25). This change has no economic interpretation, but this results in a good comparison of the fitted and actual yields. I change equation (2.25) into

$$\ln Y_t^{(\tau)} = A_\tau - B_\tau' P_t + u_t^{(\tau)}. \tag{4.3}$$

This transformation causes the yields to be reflected about the x-axis, and should thus provide a better fit to the actual yields. Figure 9 shows this transformation. We can see that the fitted yields follow the observed yields quite well. Although the transformation is not justified by any economic factors and only provides better fits for bond yields with a maturity larger than 36 months, it seems that the model is able to predict the observed yields to some extent. For $\tau = 120$ months, the largest error is 0.97%. When decreasing the time-to-maturity, the largest error keeps increasing. It should be noted that all the transformations are done ex-post, and the reflection of yields is purely based on observations.

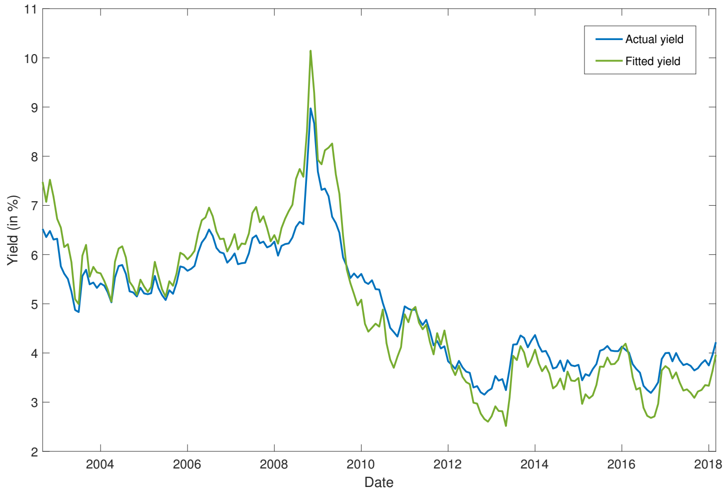

Figure 9: Fitted corporate bond yields for the $K = 5$ factor model reflected about the x-axis. This figure shows the fitted corporate bond yields of the $K = 5$ factor ACM model after it is reflected about the x-axis. The blue line corresponds to the observed yield, the green line corresponds to the fitted yield.

Table 5 shows the yield fitting errors of the reflected yields. We can see that the standard deviation is considerably lower for all maturities except $\tau = 6$ months. The skewness and kurtosis are also reduced for all maturities, except for $\tau = 60$ months. The first order autocorrelation is not affected much, but the sixth order is lower for the bonds with a maturity longer than 48 months. In general, the reflection about the x-axis is more effective as the time-to-maturity is increased, except for $\tau = 60$, where the errors have a large negative skewness and high kurtosis.

Table 5: Reflected yield fitting error properties for different time-to-maturities of corporate bonds. This table provides a summary of the fitted yield errors obtained from the five-factor ACM model with yields reflected about the x-axis. The sample period is July 2002, to March, 2018. The statistics considered are the mean, standard deviation, skewness, kurtosis, and the first and sixth order autocorrelations, $\rho(1)$ and $\rho(6)$ respectively. τ corresponds to the time-to-maturity in months.

Summary statistic	$\tau = 6$	$\tau = 24$	$\tau = 36$	$\tau = 48$	$\tau = 60$	$\tau = 120$
Yield errors ($\hat{u}$)						
Mean	0.000	0.000	0.000	0.000	0.000	0.000
Std. Dev.	4.582	2.213	1.527	1.373	1.551	0.472
Skewness	1.388	0.652	0.159	-1.101	-2.303	-0.450
Kurtosis	4.819	2.765	1.946	5.045	10.122	2.706
$\rho(1)$	0.970	0.967	0.947	0.925	0.926	0.919
$\rho(6)$	0.841	0.787	0.752	0.503	0.499	0.614

5 Conclusion

The goal of this paper was to replicate the findings of Adrian et al. (2013) and to extend their research by applying the model to corporate bonds. I start by obtaining pricing factors of U.S. Treasury zero-coupon yields, constructed using Nelson-Siegel-Svensson curves. Furthermore, I construct portfolios of U.S. corporate bonds by removing all non-investment grade bonds and bonds with a negative yield-to-maturity. After the data cleaning, I sort the bonds by maturity, and created maturity sorted portfolios at each point in time. This is done by taking the average of all yield-to-maturities of bonds maturing within six months of the specified time-to-maturity. Next, I apply the three linear regressions of the ACM model to the Treasury and corporate bond yields and test the correct amount of pricing factors. This model provides parameters to reconstruct the yield curve recursively. The model is significantly faster than traditional estimation methods, which rely on maximum likelihood procedures. The ACM model with five pricing factors is the correct specification for Treasury, and for corporate bond yields, the model with three or five factors gives a correct specification. Moreover, for Treasury yields, the predicted yields provide a very tight fit to the observed yields. The largest yield error is only 0.095%. The model also produces very accurate forecasts for the excess returns, where the largest error is only 0.30%. On the other hand, the ACM model does not predict corporate bond yields accurately. For shorter maturities, the predicted yields follow the actual yields during a stable economy, but whereas yields rise sharply in the financial crisis, the model predicts that yields go down. For longer maturities, the predicted yields move in the opposite direction as the observed yields. When reflecting the fitted yields about the x-axis, they follow the observed yields very well. For bonds with a time-to-maturity of 10 years, the largest error is only 0.97%. The yield and excess return errors of the standard model are heavily skewed, have a high kurtosis and standard deviation, which are all signs of a bad fit.

The ACM model produces biased yields for both the Treasury and corporate bond yields. For shorter time-to-maturities of Treasury yields, the bias is relatively small. However, due to the recursive method of obtaining pricing parameters A_τ and B_τ, the yields of larger time-to-maturities are too low. The fitted corporate bond excess holding returns are also much less volatile compared to the observed excess returns. The ACM model needs to be adapted to correct for this bias. For further research, it might be worth extending the ACM model by including a risk parameter that captures credit risk. Since most advanced credit risk models heavily rely on maximum likelihood, they could greatly benefit from the computationally faster model introduced by Adrian et al. (2013).

References

Abrahams, M., Adrian, T., Crump, R. K., Moench, E., & Yu, R. (2016). Decomposing real and nominal yield curves. *Journal of Monetary Economics*, *84*, 182–200. Retrieved from https://doi.org/10.1016/j.jmoneco.2016.10.006

Adrian, T., Crump, R. K., & Moench, E. (2013). Pricing the term structure with linear regressions. *Journal of Financial Economics*, *110*(1), 110–138. Retrieved from https://doi.org/10.1016/j.jfineco.2013.04.009

Anderson, T. W. (1951). Estimating linear restrictions on regression coefficients for multivariate normal distributions. *The Annals of Mathematical Statistics*, *22*(3), 327–351. Retrieved from https://doi.org/10.1214/aoms/1177729580

Ang, A., & Piazzesi, M. (2003). A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables. *Journal of Monetary Economics*, *50*(4), 745–787. Retrieved from https://doi.org/10.1016/S0304-3932(03)00032-1

Bauer, M. D., Rudebusch, G. D., & Wu, J. C. (2012). Correcting estimation bias in dynamic term structure models. *Journal of Business & Economic Statistics*, *30*(3), 454–467. Retrieved from https://doi.org/10.1080/07350015.2012.693855

Chen, R., Liu, B., & Cheng, X. (2010). Pricing the term structure of inflation risk premia: Theory and evidence from TIPS. *Journal of Empirical Finance*, *17*(4), 702–721. Retrieved from https://doi.org/10.1016/j.jempfin.2010.01.002

Cochrane, J. H., & Piazzesi, M. (2005). Bond risk premia. *American Economic Review*, *95*(1), 138–160. Retrieved from http://doi.org/10.1257/0002828053828581

Cochrane, J. H., & Piazzesi, M. (2009). *Decomposing the yield curve.* Retrieved from https://faculty.chicagobooth.edu/john.cochrane/research/Papers/interest_rate_revised.pdf (Unpublished)

De Pooter, M., Ravazzolo, F., & Van Dijk, D. J. (2010). *Term structure forecasting using macro factors and forecast combination* (FRB International Finance Discussion Paper No. 993). Federal Reserve Bank. Retrieved from https://doi.org/10.2139/ssrn.1585174

Duffee, G. R. (2002). Term premia and interest rate forecasts in affine models. *The Journal of Finance*, *57*(1), 405–443. Retrieved from https://doi.org/10.1111/1540-6261.00426

Gürkaynak, R. S., Sack, B., & Wright, J. H. (2007). The U.S. Treasury yield curve: 1961 to the present. *Journal of Monetary Economics*, *54*(8), 2291–2304. Retrieved from https://doi.org/10.1016/j.jmoneco.2007.06.029

Hamilton, J. D., & Wu, J. C. (2012). Identification and estimation of Gaussian affine term structure models. *Journal of Econometrics*, *168*(2), 315–331. Retrieved from https://doi.org/10.1016/j.jeconom.2012.01.035

Haubrich, J. G., & Dombrosky, A. M. (1996). Predicting real growth using the yield curve. *Economic Review*, *32*(1), 26–35. Retrieved from https://pdfs.semanticscholar.org/215c/c13b61f25556c25e0d3e14a76efe79b7d458.pdf

Jarrow, R. A., Lando, D., & Turnbull, S. M. (1996). A Markov model for the term structure of credit risk spreads. *The Review of Financial Studies*, *10*(2), 481–523. Retrieved from https://doi.org/10.1093/rfs/10.2.481

Joslin, S., Priebsch, M., & Singleton, K. J. (2014). Risk premiums in dynamic term structure models with unspanned macro risks. *The Journal of Finance*, *69*(3), 1197–1233. Retrieved from https://doi.org/10.1111/jofi.12131

Joslin, S., Singleton, K. J., & Zhu, H. (2011). A new perspective on gaussian dynamic term structure models. *The Review of Financial Studies*, *24*(3), 926–970. Retrieved from https://doi.org/10.1093/rfs/hhq128

Litterman, R., & Scheinkman, J. (1991). Common factors affecting bond returns. *Journal of fixed income*, *1*(1), 54–61. Retrieved from https://www.math.nyu.edu/faculty/avellane/

Litterman1991.pdf

Nelson, C. R., & Siegel, A. F. (1987). Parsimonious modeling of yield curves. *The Journal of Business*, *60*(4), 473–489. Retrieved from http://www.jstor.org/stable/2352957

Pearson, N. D., & Sun, T. (1994). Exploiting the conditional density in estimating the term structure: An application to the Cox, Ingersoll, and Ross model. *The Journal of Finance*, *49*(4), 1279–1304. Retrieved from https://doi.org/10.2307/2329186

Sheikh, A. Z., & Qiao, H. (2010). Non-normality of market returns: A framework for asset allocation decision making. *The Journal of Alternative Investments*, *12*(3), 8-35,4. Retrieved from https://doi.org/10.3905/JAI.2010.12.3.008

Svensson, L. (1994). *Estimating and Interpreting Foreward Interest Rates: Sweden 1992-1994* (Working Paper No. 4871). National Bureau of Economic Research. Retrieved from https://doi.org/10.3386/w4871

Zhou, C. (2001). The term structure of credit spreads with jump risk. *Journal of Banking & Finance*, *25*(11), 2015–2040. Retrieved from https://doi.org/10.1016/S0378-4266(00)00168-0

List of Tables

List of Figures

A Bias of Treasury yields

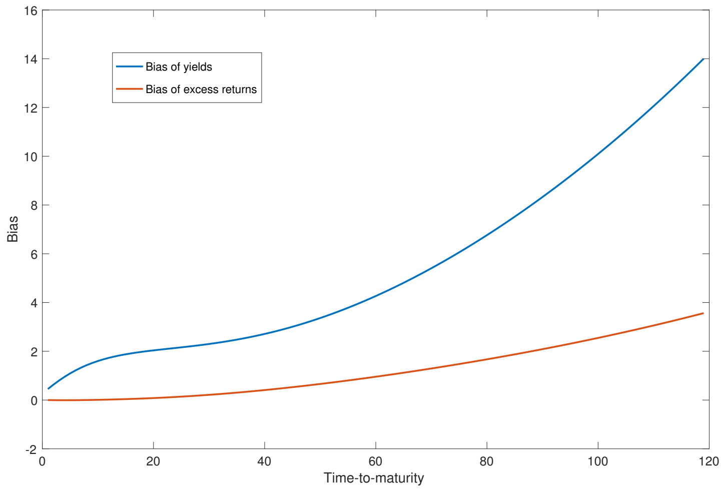

Figure 10: Bias of fitted Treasury yields and excess returns. This figure shows how the bias of the fitted Treasury yields and excess returns increases. The yields and excess returns are fitted using the ACM model with $K = 5$ factors.

B Corporate bond yield fitting with three pricing factors

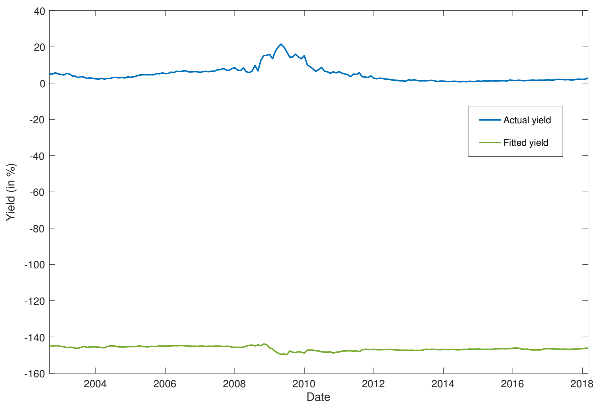

(a) $\tau = 6$

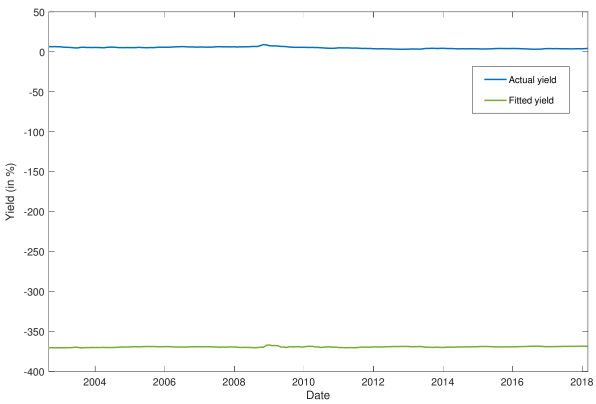

(b) $\tau = 120$

Figure 11: Fitted corporate yields for the $K = 3$ factor model for different time-to-maturities. The plots in this figure show the fitted yields of corporate bonds obtained from the three factor ACM model. The solid blue line is the observed yield, the solid green line is the fitted yield.

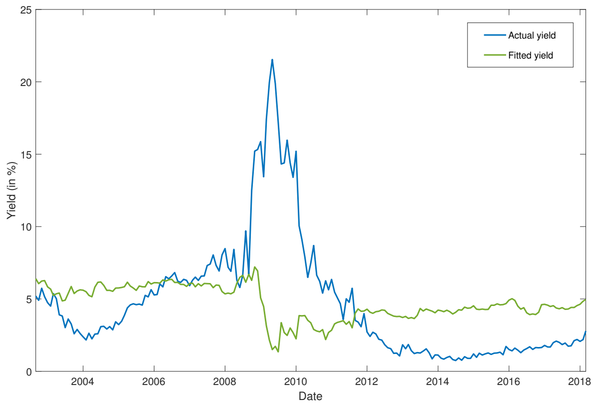

(a) Fitted yields with $\tau = 6$ months

(b) Fitted yields with $\tau = 120$ months

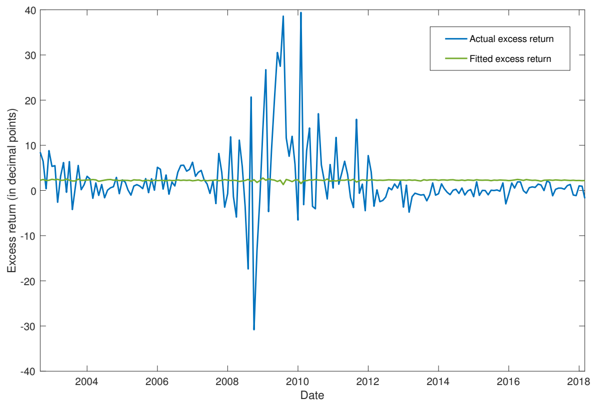

(c) Fitted excess returns with $\tau = 6$ months

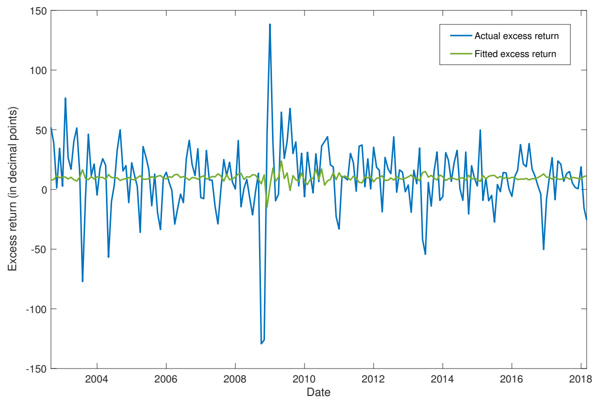

(d) Fitted excess returns with $\tau = 120$ months

Figure 12: Bias corrected fitted corporate bond yields and excess returns for the $K = 3$ factor model for different time-to-maturities. This figure presents plots of fitted yields and excess returns of the three-factor ACM model for corporate bonds after a bias correction. The solid blue line is the observed yield and excess return, the solid green line is the fitted yield and excess return.

C Code for Treasury yields

```
% Import the NSS parameters from the GSW file or load startACM.mat
load('startACM.mat')

% Calculate yields and excess returns
Yt = yields(GSW);
rx = exReturns(GSW);

% Calculate the ACM model with the wanted amount of pricing factors
K = 5;
[m, P, S, b, s, l0, l1, pf, aPF, VARV] = ACM(Yt, rx, K);

% Or calculate the CP model
% fYt = forwardRate(GSW);
% [m, P, S, b, s, l0, l1, pf, aPF, VARV] = CP(Yt, rx, fYt, rx(1:10, :));

% Perform specification tests
[anderson, pValAnderson] = canonTest(VARV, rx);
[wald, pValWald] = betaWald(b, S, s);

% Recursively calculate the predicted yields
[eYt, B] = recursivePricing(m, S, P, s, l0, l1, GSW, pf);

% Calculate the excess returns
erx = exReturnsFit(B, l0, l1, aPF, S, s, VARV);

% Select the twelve yields and correct the bias for the yields and returns
sYt = crossYields(Yt);
bYt = biasCorr(sYt, eYt);
brx = biasCorr(rx, erx);

% Calculate errors
errorYt = sYt - bYt;
errorRX = rx - brx;
```

D Code for corporate bond yields

```
% Import the bond data from the WRDS file or load startExt.mat, calculate yields and
% excess returns. The risk-free rate price is from the Treasury yield model
% fullYt = createYields(bondpf);
load('startExt.mat');
rx = exReturnsExt(fullYt, riskfreeRatePrice);

% Select the cross-section of 40 and 12 yields
Yt = selectYields(fullYt);
sYt = crossYieldsExt(Yt);

% Calculate the ACM model with the wanted amount of pricing factors
K = 5;
[m, P, S, b, s, l0, l1, pf, aPF, VARV] = ACM(Yt, rx, K);

% Perform specification tests
[anderson, pValAnderson] = canonTest(VARV, rx);
[wald, pValWald] = betaWald(b, S, s);

```

```
% Recursively calculate the predicted yields
[eYt, B] = recursivePricingExt(m, S, P, s, l0, l1, pf, riskfreeRatePrice);

% Calculate the excess returns
erx = exReturnsFit(B, l0, l1, aPF, S, s, VARV);

% Correct the bias for the yields and returns
bYt = biasCorr(sYt, eYt);
brx = biasCorr(rx, erx);

% Obtain yields recursively and calculate log price by subtracting instead
% of summing
[cYt, Bc] = recursivePricingExtChanged(m, S, P, s, l0, l1, pf, riskfreeRatePrice);

% Calculate the excess returns
crx = exReturnsFit(Bc, l0, l1, aPF, S, s, VARV);

% Correct the bias for the yields and returns
cbYt = biasCorr(sYt, cYt);
cbrx = biasCorr(rx, crx);

% Calculate errors
errorYt = sYt - bYt;
errorRX = rx - brx;

errorcYt = sYt - cbYt;
errorcRX = rx - cbrx;
```