Expected Returns in Treasury Bonds

Jake Mc Leroth

and

Nicole Lyne

Research assignment presented in the partial fulfilment
of the requirement for the degree of
BComHons (Financial Risk Management)
at the University of Stellenbosch

Supervisor: Mr. O. Sihlobo

Degree of confidentiality: A October 2024

PLAGIARISM DECLARATION

1. Plagiarism is the use of ideas, material and other intellectual property of another's work and to present it as our own.
2. We agree that plagiarism is a punishable offence because it constitutes theft.
3. Accordingly, all quotations and contributions from any source whatsoever (including the internet) have been cited fully. We understand that the reproduction of text without quotation marks (even when the source is cited) is plagiarism.
4. We also understand that direct translations are plagiarism.
5. We declare that the work contained in this research assignment, except otherwise stated, is our original work and that we have not previously (in its entirety or in part) submitted it for grading in this research assignment or another research assignment.

24913693	
Student number	Signature
J. Mc Leroth	25 October 2024
Initials and surname	Date
24744433	
Student number	Signature
N. Lyne	25 October 2024
Initials and surname	Date

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1

INTRODUCTION

In this paper, we explore the framework for the prediction of bond returns provided by Cieslak and Povala (2011) in the context of the United States (US) and South Africa (SA). We make adjustments to this framework to allow for real-world application with bonds traded in the market. We also introduce a new idea for the model to allow for better results in the South African context. Furthermore, we evaluate the performance of these models, both in-sample and out-of-sample by looking at various performance measures.

CHAPTER 2

LITERATURE REVIEW

2.1 INTRODUCTION

The study of the drivers of excess bond returns and the methods used to predict them is an important area of financial research. Cochrane and Piazzesi (2005) seminal paper laid the groundwork for understanding the predictability of excess bond returns, using a single return forecasting factor, which has since been extended by further research. Cieslak and Povala (2011) and Adrian and Moench (2013) (ACM) expanded on this foundational work to explain time-varying term premia in bond markets. Their predictability performance is measured through metrics such as adjusted R^2 (this will be referred to as R^2).

While the literature focusses on the predictability of bond returns, the practical application of these models for investment decisions has been overlooked. Moreover, the available literature is based on US data, which raises questions about the robustness and applicability of these models in a South African context. Methodologies and key findings of bond return predictability models will be discussed, as well as methods of evaluating investment performance (Brooks, 2021).

2.2 UNDERSTANDING BOND YIELDS AND THEIR DETERMINANTS

2.2.1 Key Components of Bond Yields

The yield on a bond depends on the current short-term interest rate, expected future short-term interest rates, and the term premium, which is the difference between the yield on the bond and what the market expects interest rates to be in the future. In other words, the term premium measures the excess return investors require for holding a long-term bond to maturity rather than holding a series of short-maturity bonds (Brooks, 2021).

The yield curve, which is made up of bond yields of different maturities, also plays a crucial role in understanding bond yields. An upward-sloping yield curve suggests investors expect interest rates to rise and term premia to be positive (and increase with maturity). Conversely, a downward-sloping yield curve might indicate expectations of future interest rate cuts, which is a potential indication of a recession (Brooks, 2021).

2.2.2 The Expectations Hypothesis (EH)

To have a deeper understanding of the policy and non-policy drivers of bond yields, it is essential to discuss the Expectations Hypothesis (EH). The EH suggests that bond yields are affected by current and expected future interest rates only and changes in these interest rates have an effect on the slope and level of the yield curve. The EH does not account for time-varying term premia as it assumes that term premia are constant over time, although they may differ by maturity. When central banks increase the interest rate, expected future interest rates for the next few years also rise, but expected future interest rates for years further into the future change by less. Since the yield on a bond is the average of the interest rates over its lifetime, long-maturity bond yields will rise less than short-maturity bond yields. This causes the curve to flatten (Brooks, 2021).

2.2.3 Monetary Policy and Yield Curve Dynamics

It is also important to understand what drives interest rates. Short-term interest rates are controlled by central banks and their behaviour can be described by the monetary policy equation, known as the "Taylor Rule", which describes how the central bank sets the interest rate based on short-term forecasts of unemployment and inflation (Taylor, 1993). Central banks set a low interest rate when inflation is below their target level or when output is lower than their full employment level. The changes in these interest rates will affect longer-maturity bond yields (Brooks, 2021).

2.2.4 Impact of Macroeconomic Factors on Bond Yields

Monetary policy is not the only variable influencing long-maturity bond yields. The level and shape of the yield curve are also affected by non-monetary policy factors. Variations in trend growth and long-term inflation expectations, along with changes in inflation uncertainty and short-term fluctuations in the business cycle, as well as shifting demand for highly liquid, safe-haven assets, significantly impact bond yields (Brooks, 2021).

2.3 FRAMEWORKS FOR BOND RETURN PREDICTABILITY

2.3.1 Cochrane and Piazzesi

Cochrane and Piazzesi (2005) challenge the EH by demonstrating the importance of time-varying term premia in US government bonds. They run regressions of one-year excess returns on initial forward rates and find that excess returns on one- to five-year maturity bonds can be predicted by a single return forecasting factor, a linear combination of forward rates, achieving an in-sample R^2 of up to 0.44. This forecasting factor embeds many widely used predictors of bond returns and effectively captures the time variation in the expected return of all bonds.

2.3.2 Cieslak and Povala

Cieslak and Povala (2011) build upon and extend the ideas introduced by Cochrane and Piazzesi (2005). They further investigate how macroeconomic variables, including inflation and the yield curve, impact bond risk premia and returns. Their work incorporates the insights from Cochrane and Piazzesi (2005), adding more depth to the understanding of bond risk premia and improving predictive accuracy by integrating additional economic factors.

The authors first decompose the yield curve over time into the sum of short rate expectations and the risk premium for a specific maturity. The short rate is then made up of the sum of the real rate and the expected future inflation. This is known as the Fisher equation (Fisher, 1907). Inflation itself is also broken down into a persistent (long-run mean inflation) and a transitory component. This leaves a final decomposition where the yield is broken down into a persistent component that captures smooth adjustments in short rate expectations and a transitory component that contains the transitory part of short rate expectations as well as the risk premium. This transitory part is called “cycles”. This decomposition shows that the yield curve is affected by three different economic frequencies: generational frequency connected to persistent inflation expectations, business cycle frequency related to transitory short rate expectations and term premia frequency. Additionally, this decomposition furthers our understanding of the link between term premia, macroeconomic conditions and yield curve factors. If the long-run Fisher effect is true, yields and persistent inflation expectations should be cointegrated, meaning that cycles should possess predictive power of bond returns.

The authors use a exponentially weighted moving average of past core CPI inflation as a proxy for persistent inflation (trend inflation) in the yield decomposition. Running a linear regression between zero-coupon yields and the trend inflation results in residuals, which represent the cycles, that are stationary (ie, mean-reverting). This is shown by the augmented Dickey-Fuller test. Because of this stationarity, the cycles can be used for prediction of bond returns. To do this, the authors first get the cycles across all maturities. The first prediction is between the return for a specific maturity and the cycles across all maturities (1, 2, 5, 7, 10, 20 years). It was found that the cycles forecast a large portion of variation in excess bond returns and the regression produced R^2 values from 42% to 57% across maturities. This is a significant improvement compared to the predictability achieved by Cochrane and Piazzesi (2005) with forward rates.

The second predictive regression was between the average return and the cycles. The cycles were split into one-year maturity and then the average of the remaining cycles. The reason for this split is that the cycles over one year capture the variation in short rate expectations and do not reflect the risk premia as the risk premia for one year is zero. The average of the rest of the maturities then captures the variation in risk premia as the impact of the short rate expectation subsides as maturity increases. The fitted values from this model are what the authors call the return predicting factor. This return predicting factor is then used in a regression for the return of a specific maturity. This results in R^2 values from 0.38 to 0.55 across maturities. Essentially, the authors are capturing information from all maturities into one factor, which is then mapped back to specific maturities. This is a powerful tool as this information incorporated into a single factor only performs slightly worse than the original, more complex regression. This demonstrates the robustness and utility of the return predicting factor as a powerful tool for forecasting bond excess returns.

2.3.3 Adrian, Crump, and Moench

Adrian and Moench (2013) also challenge the EH by demonstrating time-varying risk premia in the U.S. Treasury bond market. They introduce a method to estimate bond risk premia and predict excess bond returns, using five principal components of yields as factors. Their model outperforms Cochrane and Piazzesi's model in out-of-sample predictability, showing improvements in R^2 values. The authors acknowledge the drawbacks of affine models and their subsequent maximum likelihood estimation of parameters and instead propose a regression-based approach to pricing interest rates. First, they decompose observable pricing factors into their predictable components and factor innovations by performing a regression of the factors on their lagged values. Then, they estimate Treasury return exposures with respect to the lagged values of these pricing factors and the current innovations in the pricing factors. Finally, they determine the market price of risk parameters by conducting a cross-sectional regression, where they relate the exposures of returns to the lagged pricing factors with the exposures to contemporaneous innovations in the pricing factors. The authors find that the first three principal components of treasury yields are insufficient to capture the cross-section of treasury returns; hence, they use the first five principal components as pricing factors.

2.4 CONCLUSION

The models explored in the literature focus on bond return predictability by directly challenging the EH. The applicability of this research in an investment setting has not yet been explored, which poses a question as to whether these models can be used to make effective investment decisions. Furthermore, it is crucial to examine whether the findings and models that have been evaluated using US data can be effectively translated to the South African bond market, given its unique economic and financial dynamics.

CHAPTER 3

DATA

All data series are obtained from the Bloomberg database.

3.1 INFLATION

The monthly US and SA CPI for all Urban Consumers Less Food & Energy (Core CPI) is used to construct the persistent inflation variable τ_t. This CPI data represents the average price an urban consumer pays for a basket of goods, excluding food and energy. Food and energy are excluded because these prices can be very volatile (Federal Reserve Bank of St. Louis, 2024). CPI is measured by the Bureau of Labour Statistics in the US and STATS SA in SA. Each dataset is downloaded from Bloomberg and the US data runs from 31/12/1966 to 29/02/2024, while the SA data runs from 29/02/1980 to 31/05/2024. Core CPI inflation is defined as the year-on-year simple growth rate in the core CPI index.

Inflation data for a certain month are only released in the middle of the next month. To deal with this lag when estimating the persistent inflation component, the CPI data that are available at the end of the previous month is used in the current month. For example, the calculation of the persistent inflation component for January 2024 uses inflation data up until December 2023.

A proxy for the long-run mean inflation (trend inflation) is created, τ_t^{proxy}, and it is an exponentially weighted moving average of the past 10 years of monthly core CPI:

$$\tau_t^{proxy} = \frac{\sum_{i=0}^{119} v^i CPI_{t-i}}{\sum_{i=0}^{119} v^i} \tag{3.1}$$

where v is the discounting factor (gain parameter). We take $v = 0.987$ (Cieslak and Povala, 2011).

To visually see the smoothing done on inflation, Figure 3.1 and Figure 3.2 show the year-on-year inflation as well as the proxy for trend inflation, τ_t^{proxy}, in the US and SA.

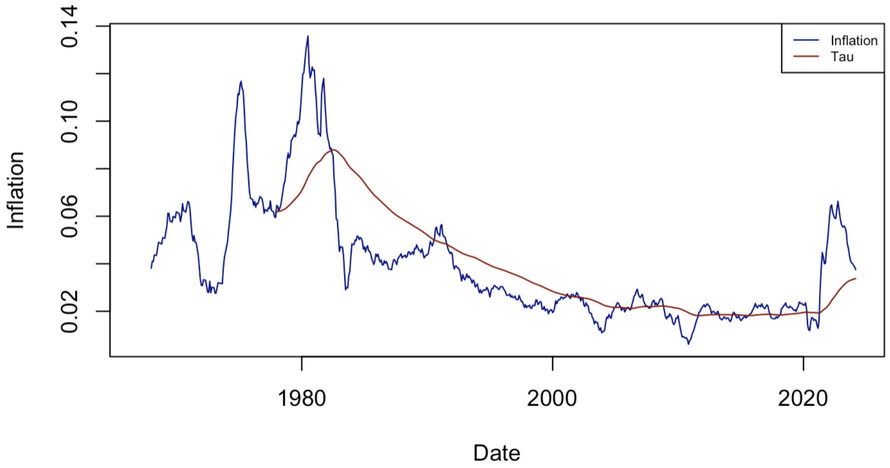

Figure 3.1: A time series of US inflation and of the proxy for the US trend inflation

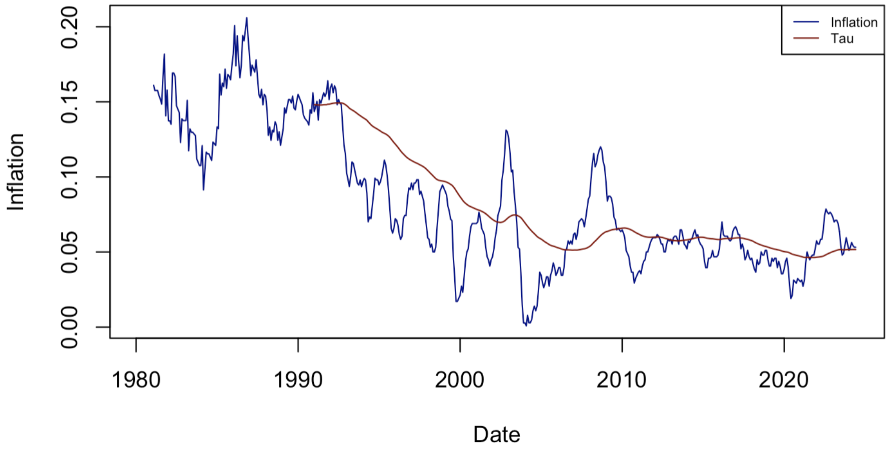

Figure 3.2: A time series of SA inflation and of the proxy for the SA trend inflation

3.2 ZERO-COUPON YIELDS

Monthly US and SA zero-coupon yields for bonds with 1-, 2-, 5-, 7-, 10- and 20-year maturities are obtained from Bloomberg. Zero-coupon bonds differ from regular coupon-paying bonds in that they do not make periodic interest payments. The zero-coupon yields reflect the yield to maturity

(YTM) of these hypothetical bonds across various maturities. The SA data ranges from 1995 to 2024, while the US data ranges from 1969 to 2024.

Figure 3.3 and Figure 3.4 show the yields for zero-coupon bonds with different maturities, as well as trend inflation, τ_t^{proxy}, in US and SA respectively.

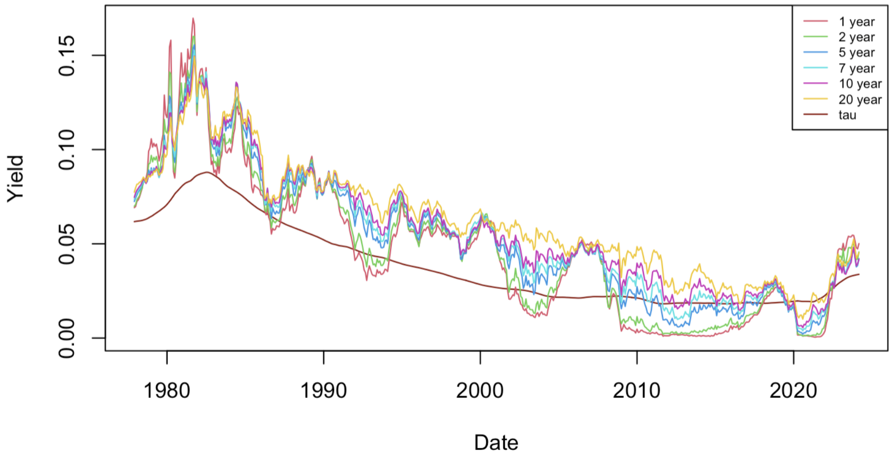

Figure 3.3: US yields of zero-coupon with different maturities as well as trend inflation

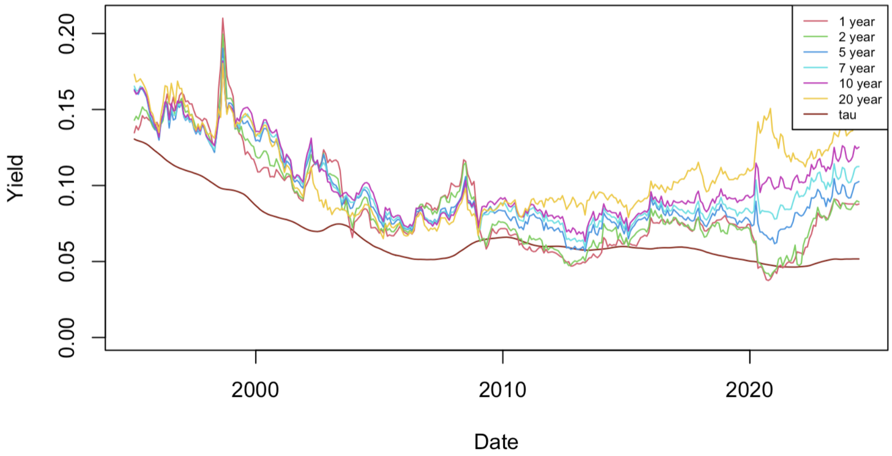

Figure 3.4: SA yields of zero-coupon with different maturities as well as trend inflation

3.3 TRADED BONDS

Data for traded bonds across different maturities for the US and SA were obtained from Bloomberg. The yields are depicted in Figure 3.5 and Figure 3.6. The SA data ranges from 1995 to 2024 and the US data ranges from 1962 to 2024.

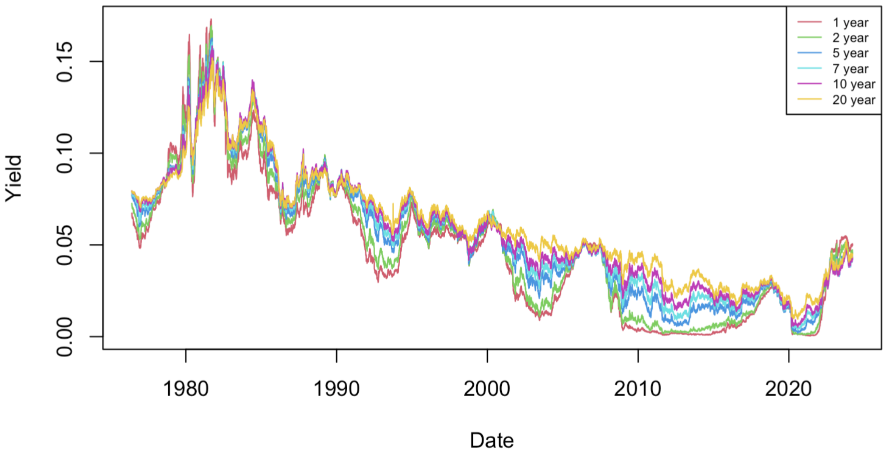

Figure 3.5: US yields of traded bonds with different maturities

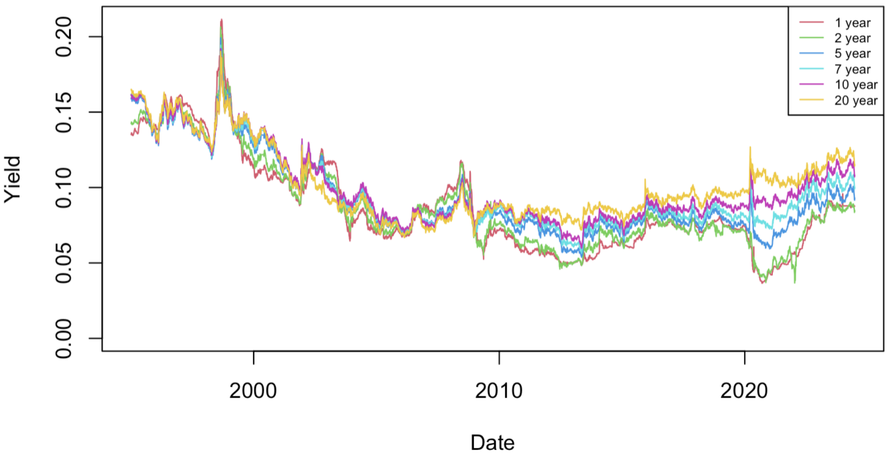

Figure 3.6: SA yields of traded bonds with different maturities

CHAPTER 4

METHODOLOGY

4.1 YIELDS AND CYCLES

Borrowing notation from Cieslak and Povala (2011), the yield at time t of a zero-coupon bond with n years to maturity is denoted as $y_t^{(n)}$. As mentioned earlier, the yield is decomposed into the following equation:

$$y_t^{(n)} = b_0^{(n)} + b_\tau^{(n)}\tau_t + b_r^{(n)} r_t + b_\pi^{(n)}\pi_t^c + rpy_t^{(n)} \tag{4.1}$$

where τ_t denotes trend inflation, π_t^c transitory inflation, r_t the real rate and $rpy_t^{(n)}$ the risk premium. The transitory part of this equation, consisting of π_t^c, r_t and $rpy_t^{(n)}$, can be combined and labelled as cycles. This is denoted as $\tilde{c}_t^{(n)}$. Hence, it follows:

$$y_t^{(n)} = b_0^{(n)} + b_\tau^{(n)}\tau_t + \tilde{c}_t^{(n)} \tag{4.2}$$

$$\tilde{c}_t^{(n)} = b_r^{(n)} r_t + b_{\pi^c}^{(n)}\pi_t^c + rpy_t^{(n)} \tag{4.3}$$

The cycles represent the variation in the short rate expectations and risk premia over time. According to the long-run Fisher effect, $y_t^{(n)}$ and τ_t are cointegrated. Hence, the cycles are stationary and should provide predictability in yield movements and bond returns. We obtain the coefficients and cycles in Equation 4.2 by running an Ordinary Least Squares (OLS) regression:

$$y_t^{(n)} = b_0^{(n)} + b_\tau^{(n)}\tau_t^{proxy} + \epsilon_t^{(n)} \tag{4.4}$$

The residuals obtained from this regression correspond to the observed cycles ($\epsilon_t^{(n)} = c_t^{(n)}$). This regression is done to obtain the cycles for the 1- to 20-year zero-coupon bonds. The τ_t^{proxy} is as defined in Equation 3.1. The yields, trend inflation and cycles are then tested for stationarity using the augmented Dickey-Fuller test to check for a cointegrating relationship between yields and trend inflation. This process is done on both US and SA data.

4.2 BOND RETURNS AND PREDICTION

The cycles contain information about the short rate and bond risk premia. Hence, Cieslak and Povala (2011) use the cycles obtained from Equation 4.4 to forecast future returns by using a regression model:

$$rx_{t+12}^{(n)} = \delta_0 + \sum_i \delta_i c_t^{(i)} + \varepsilon_{t+12}^{(n)} \tag{4.5}$$

where $i = \{1, 2, 5, 7, 10, 20\}$ and $rx_{t+12}^{(n)} = p_{t+12}^{(n-1)} - p_t^{(n)} - y_t^{(1)}$ is the excess log return, with $p_t^{(n)}$ the log price of a zero-coupon bond with n years to maturity at time t. Since monthly data is used, $t + 12$ in $rx_{t+12}^{(n)}$ represents 12 months into the future, which is equivalent to one year ahead. This model is used to compare against the return forecasting factor which they construct, $\widehat{cf}_t$:

$$\overline{rx}_{t+12} = \gamma_0 + \gamma_1 c_t^{(1)} + \gamma_2 \bar{c}_t + \varepsilon_{t+12} \tag{4.6}$$

$$\widehat{cf}_t = \hat{\gamma}_0 + \hat{\gamma}_1 c_t^{(1)} + \hat{\gamma}_2 \bar{c}_t \tag{4.7}$$

where $\overline{rx}_{t+1}$ is the average excess return across 2- to 20-years and $\bar{c}_t = \frac{1}{5}\sum_j c_t^{(j)}$ where $j = \{2, 5, 7, 10, 20\}$. The underlying intuition is that the 1-year cycles represent the variation in short rate expectations and the average of the rest of the cycles represent the variation in risk premia across all maturities. This factor is then used in the following regression:

$$rx_{t+12}^{(n)} = \beta_0^{(n)} + \beta_1^{(n)} \widehat{cf}_t + \varepsilon_{t+12}^{(n)}. \tag{4.8}$$

The shortfall of this methodology is that the returns are calculated using zero-coupon bonds, which are not actively traded in the market. To make these models more applicable in the real world, we use yields from bonds that are traded in the market. To calculate the return on these bonds, denoted as $rp_{t+12}^{(n)}$, we use a Taylor series expansion, purposefully ignoring deterministic accrued interest to focus only on modelling uncertainty. The return, $rp_{t+12}^{(n)}$, is defined as follows:

$$rp_{t+12}^{(n)} = -(\text{MODIFIED DURATION}_t)(y_{t+12}^{*(n)} - y_t^{*(n)}) + \frac{1}{2}(\text{CONVEXITY}_t)(y_{t+12}^{*(n)} - y_t^{*(n)})^2 \tag{4.9}$$

where $y_t^{*(n)}$ is the yield of a traded bond with maturity n at time t.

4.2.1 Full Model

We then run the following regression:

$$rp_{t+12}^{(n)} = \delta_0 + \sum_i \delta_i c_t^{(i)} + \varepsilon_{t+12}^{(n)} \tag{4.10}$$

In addition to OLS regression, we apply Elastic Net regression to the full model. A challenge of multiple linear regression is that predictor variables can be highly correlated with each other. This creates multicollinearity, which makes it challenging to determine the independent effects of the predictor variables on a response variable. OLS regression does not impose any penalties for multicollinearity (Sztepanacz and Houle, 2024). Conversely, Elastic Net regression effectively deals with multicollinearity by imposing a penalty that shrinks some coefficients to zero (Lasso) and stabilises the coefficient estimates by reducing the magnitude of the coefficients towards each other, enabling them to borrow strength from each other (Ridge) (Friedman *et al.*, 2010).

4.2.2 Factor Model

The return forecasting factor is obtained by regressing the average return of traded bonds across all maturities, $\overline{rp}_{t+12}$, onto the transitory short rate expectations factor, $c_t^{(1)}$, and the average cycle, $\bar{c}_t$, using OLS regression:

$$\overline{rp}_{t+12} = \gamma_0 + \gamma_1 c_t^{(1)} + \gamma_2 \bar{c}_t + \varepsilon_{t+12} \tag{4.11}$$

The fitted values from the above regression represent the return forecasting factor, $\widehat{cf}_t$, which is then used as a return predictor across all maturities:

$$\widehat{cf}_t = \hat{\gamma}_0 + \hat{\gamma}_1 c_t^{(1)} + \hat{\gamma}_2 \bar{c}_t \tag{4.12}$$

$$rp_{t+12}^{(n)} = \beta_0^{(n)} + \beta_1^{(n)} \widehat{cf}_t + \varepsilon_{t+12}^{(n)} \tag{4.13}$$

This too is done on both US and SA data.

4.3 COMBINING US AND SA CYCLES

Due to the prominence of the US economy and the fact that South Africa is an emerging market, what occurs in the US has a strong impact on the South African bond market (Fedderke, 2021). With this in mind, we construct another regression in an attempt to improve the explainability and predictability of returns in South Africa:

$$rp_{t+12}^{(n)} = \delta_0 + \sum_i \delta_i c_t^{(i)} + \sum_i \delta_i^* c_t^{*(i)} + \varepsilon_{t+12}^{(n)} \tag{4.14}$$

Where $rp_{t+12}^{(n)}$ are the SA traded bond returns, $c_t^{(i)}$ are the SA cycles and $c_t^{*(i)}$ are the US cycles. The SA and US cycles obtained for this regression are from the same time frame of data.

4.4 OUT-OF-SAMPLE TESTING

The training set consists of at least 120 observations for all models, with the test set being the next twelve observations. Twelve predictions are made and compared to the test set. The test set is then added to the training set, and a new test set is created with the next twelve observations. This process is repeated through time, resulting in a set of predicted and observed returns. The R^2s are calculated on these out-of-sample predictions to get the out-of-sample R^2s. We also calculate the accuracy, which is the ratio of correct predictions, as follows:

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}} \tag{4.15}$$

TP (true positives) is the number of positive returns that were predicted as positive, FP (false positives) is the number of negative returns that were predicted as positive, TN (true negative) is the number of negative returns that were predicted as negative and FN (false negative) is the number of positive returns that were predicted as negative. This measure indicates what proportion of time the model was able to correctly predict the direction of the returns (Géron, 2019).

CHAPTER 5

RESULTS

5.1 US CONTEXT

5.1.1 Cycles

Figure 5.1 shows a time series of US inflation, the zero-coupon bond yields of 1- to 20-year maturities, the trend inflation (the variable τ_t^{proxy}) and the fitted yields from Equation 4.4. The cycles are represented by the differences between the smoothed lines (predicted yields) and the jagged lines (actual yields) in Figure 5.1. These cycles are shown in Figure 5.2.

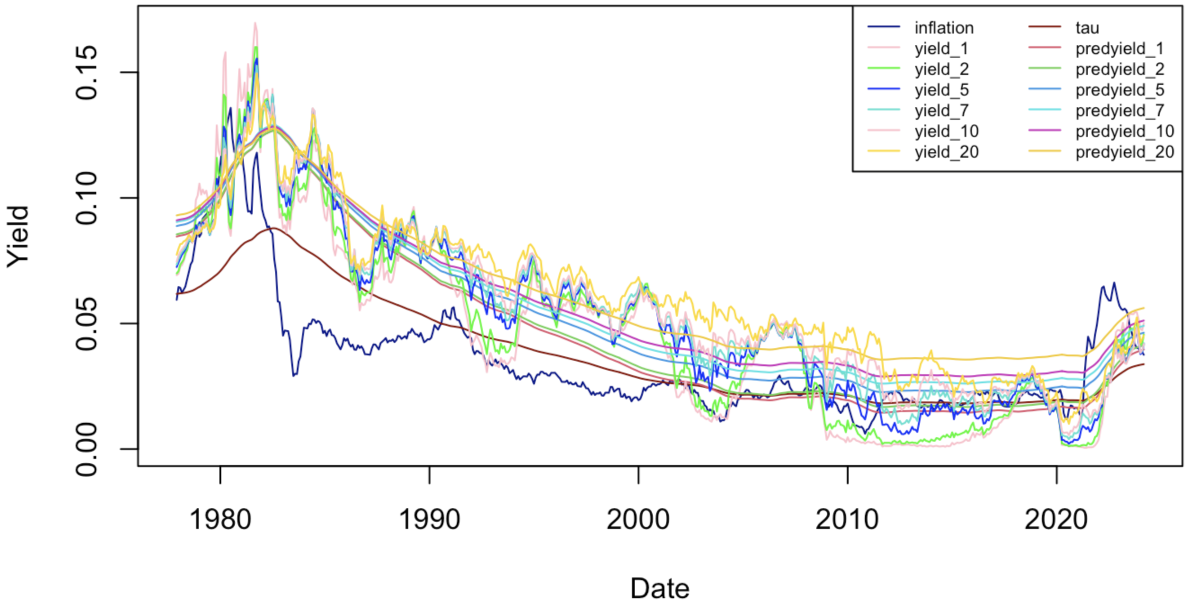

Figure 5.1: A time series of US inflation, of the US yields of maturities from 1- to 20-years, of the trend inflation (the variable τ_t^{proxy}) and the fitted yields from Equation 4.4

Figure 5.2 shows the variation in the transitory component of Equation 4.1 across maturities. The cycles for each maturity show mean reversion: when the cycles increase, they tend to come back down. This mean reversion will also be evaluated by looking at the stationarity of these cycles (residuals) (Dias and Marques, 2010). When the cycles are high, the zero-coupon yields are much higher than trend inflation and an investor might believe that these cycles (the gap between yields and inflation) will decrease due to mean reversion. This would be a good time to buy bonds, since it is expected that yields will decline in the future.

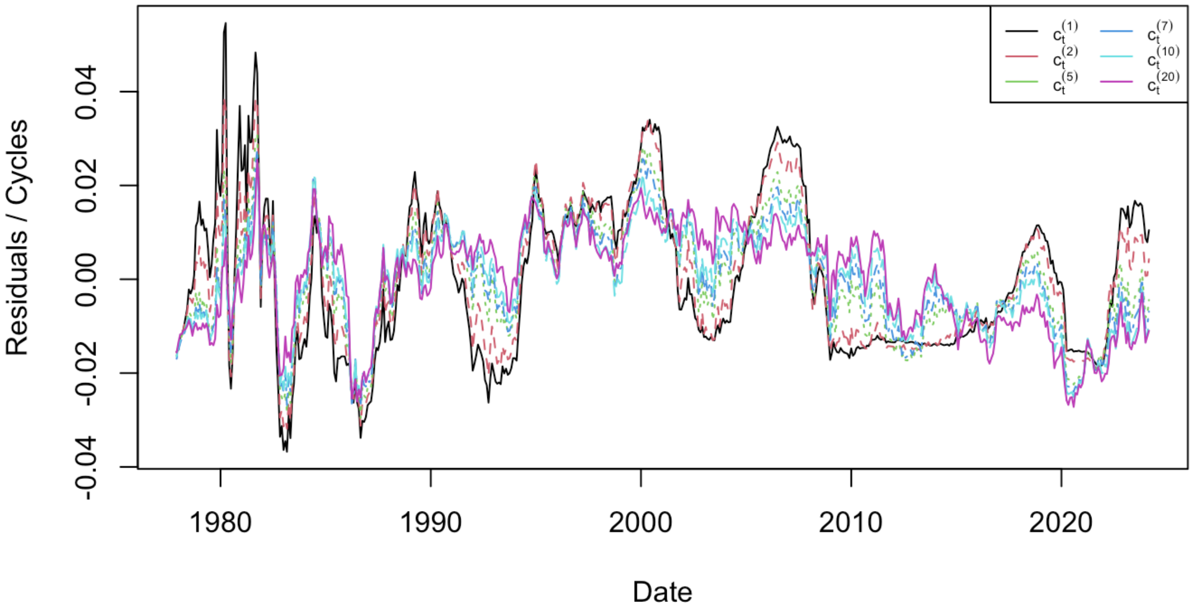

Figure 5.2: US residuals/cycles, $c_t^{(n)}$ or $\varepsilon_t^{(n)}$, as in Equation 4.4

.

5.1.2 Cointegration

Cointegration describes the situation where two variables are individually non-stationary but there exists a linear combination of the two variables that is stationary. This definition has strong economic value as the variables can be seen to have a stable relationship, where any deviations from a long-run equilibrium are temporary and are corrected over time (Stigler, 2020). As mentioned before, according to the long-run Fisher effect, yields and trend inflation should be cointegrated.

An augmented Dickey-Fuller (ADF) test is a statistical test used to check for a unit root in a time series. If a unit root exists, the time series is non-stationary. However, if there is no unit root present, the ADF test rejects the null hypothesis of a unit root, suggesting that the time series is stationary. The test was performed on the tau and yield values as well as on the cycles.

ADF Test for τ_t^{proxy} and Zero-Coupon Yields

Table 5.1 reports the p-values from the ADF test for τ_t^{proxy} and yields. The ADF tests on the yields reveal p-values larger than 0.05 at lag 1, therefore we cannot reject the null hypothesis and we conclude that the yields have a unit root and are non-stationary. We see that τ_t^{proxy} has a p-value less than 0.05 and we therefore should reject the null hypothesis. However, we can clearly see from Figure 3.1 that τ_t^{proxy} is not stationary, hence we do the ADF test at lag 2 and see that the p-value is above 0.05. Hence, we fail to reject the null hypothesis and conclude that τ_t^{proxy} is non-stationary.

Table 5.1: p-values from ADF test for τ_t^{proxy} and zero-coupon yields

No. of lags	τ_t^{proxy}	$y_t^{(1)}$	$y_t^{(2)}$	$y_t^{(5)}$	$y_t^{(7)}$	$y_t^{(10)}$	$y_t^{(20)}$
1	0.0196	0.1591	0.1869	0.1074	0.1093	0.0936	0.05446
2	0.5088	0.3659	0.3811	0.2529	0.2178	0.1857	0.1087

ADF Test for Residuals

Table 5.2 reports the p-values from the ADF test for the cycles, $c_t^{(n)}$, which are obtained from the regression in Equation 4.4. The p-values for the cycles are smaller than 0.05 at lag 1, therefore we reject the null hypothesis of a unit root and we conclude that the cycles are stationary. Since the yields and the trend inflation are non-stationary and the cycles, which are a linear combination of the yields and the trend inflation, are stationary, the yields and the trend inflation are cointegrated at all maturities. Hence, the level of the mean-reverting cycles provides insight into where yields, and therefore returns, will likely be in the future. This also means the long-run Fisher effect holds in the US market.

Table 5.2: p-values from ADF test for cointegrating residual

No. of lags	$c_t^{(1)}$	$c_t^{(2)}$	$c_t^{(5)}$	$c_t^{(7)}$	$c_t^{(10)}$	$c_t^{(20)}$
1	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01

5.1.3 Bond Returns and Predictability

The R^2 value measures how well the independent variable explains the variation in the dependent variable. A high R^2 value suggests that a large proportion of the variation in the dependent variable is explained by the independent variable (Romeo, 2020).

In a linear regression, the R^2 value always increases or at least stays the same as the number of independent variables added to the model increases. This leads to an overestimation of the fit of the model. A regression model containing more independent variables than another model may look as though it has a better fit solely because it has more variables, which can be misleading.

A solution to this issue is the adjusted R^2, which penalises the R^2 for the amount of predictor variables the model has. This makes models with a different amount of predictors comparable (James *et al.*, 2013).

Full Model

Recall the full model from Equation 4.10:

$$rp_{t+12}^{(n)} = \delta_0 + \sum_i \delta_i c_t^{(i)} + \varepsilon_{t+12}^{(n)}$$

We obtained fitted values for each bond maturity. We applied two different regression techniques, namely OLS regression and Elastic Net regression, to the full model to determine which is best for bond return predictability. In this scenario, the adjusted R^2 value measures how well the fitted values from the regression explain the variation in the returns for a traded bond with maturity n.

Regression Coefficients Studying Table 5.3, we see that, as expected, some of the regression coefficient estimates have been shrunk to zero when using Elastic Net regression. This provides a simplified model, showing only the main drivers of the bond returns.

Table 5.3: Comparison of OLS and Elastic Net regression coefficients for different horizons in the US context - full model

	OLS						Elastic Net					
	$rp^{(1)}$	$rp^{(2)}$	$rp^{(5)}$	$rp^{(7)}$	$rp^{(10)}$	$rp^{(20)}$	$rp^{(1)}$	$rp^{(2)}$	$rp^{(5)}$	$rp^{(7)}$	$rp^{(10)}$	$rp^{(20)}$
δ_0	0.0003	0.0006	0.0022	0.0037	0.0052	0.0065	0.0003	0.0006	0.0022	0.0039	0.0056	0.0065
δ_1	0.3598	-0.1307	-1.1635	-1.8930	-2.8227	-3.3574	0.3035	-0.1404	-0.9802	-1.3589	-1.7708	-2.8380
δ_2	-0.4457	0.2507	0.4571	0.9875	1.9313	3.3549	-0.4287	0.0000	0.0000	0.0000	0.0000	0.0000
δ_5	-0.5319	-1.2611	-1.7426	-5.3585	-6.5085	-11.4823	0.0000	0.0000	-1.1459	0.0000	0.0000	0.0000
δ_7	0.7916	1.3341	3.4168	8.3347	7.8368	11.7014	0.0000	0.1309	3.2301	4.5069	4.3347	1.8283
δ_{10}	0.7372	1.7208	4.2516	6.0248	9.6870	9.6430	0.9929	1.7827	4.0740	-0.1906	1.1317	9.7083
δ_{20}	-0.4898	-1.0976	-3.1608	-5.0063	-6.2456	-5.4113	-0.4475	-0.9773	-3.1338	0.0000	0.0000	-4.4449

In-Sample adjusted R^2 Using OLS and Elastic Net Regression Studying Table 5.4, the adjusted R^2s using OLS regression range from 25.55% to 39.28%, which shows that up to 39.28% of the variation in traded bond returns can be explained by this model. Moreover, it is clear to observe that the adjusted R^2s from the OLS regression are higher than the adjusted R^2s obtained from the Elastic Net regression. It should be noted that this will always be true in-sample since OLS regression explicitly minimises the residual sum of squares (James *et al.*, 2013).

Table 5.4: In-sample adjusted R^2 values for different maturities using OLS and Elastic Net regression in US context - full model

adjusted R^2	$rp^{(1)}$	$rp^{(2)}$	$rp^{(5)}$	$rp^{(7)}$	$rp^{(10)}$	$rp^{(20)}$
OLS	0.2555	0.2883	0.3530	0.3833	0.3928	0.3900
Elastic Net	0.2543	0.2868	0.3528	0.3528	0.3611	0.3823

Factor Model

Recall the OLS regression from Equation 4.11:

$$\overline{rp}_{t+12} = \gamma_0 + \gamma_1 c_t^{(1)} + \gamma_2 \bar{c}_t + \varepsilon_{t+12}$$

In this case, the adjusted R^2 value measures how well the fitted values from the regression, $\widehat{cf}_t$, explain the variation in the average return. From this regression, we achieved an adjusted R^2 of 0.3434. This indicates that approximately 34.34% of the variation in the average returns across different maturities is explained by the variation in short rate expectations, $c_t^{(1)}$, and the variation in risk premia across all maturities, $\bar{c}_t$.

The fitted values from this regression correspond to the single return forecasting factor, $\widehat{cf}_t$, as shown in Equation 4.12:

$$\widehat{cf}_t = \hat{\gamma}_0 + \hat{\gamma}_1 c_t^{(1)} + \hat{\gamma}_2 \bar{c}_t$$

Which we regress in Equation 4.13:

$$rp_{t+12}^{(n)} = \beta_0^{(n)} + \beta_1^{(n)} \widehat{cf}_t + \varepsilon_{t+12}^{(n)}$$

Fitted Values & Regression Coefficients Figure 5.3 shows a plot of the return forecasting factor, $\widehat{cf}_t$, over time. The higher the return forecasting factor, the more likely it is that bonds will rally going forward. The lower the factor is, the more likely it is that bonds will lose value going forward (Cochrane and Piazzesi, 2005).

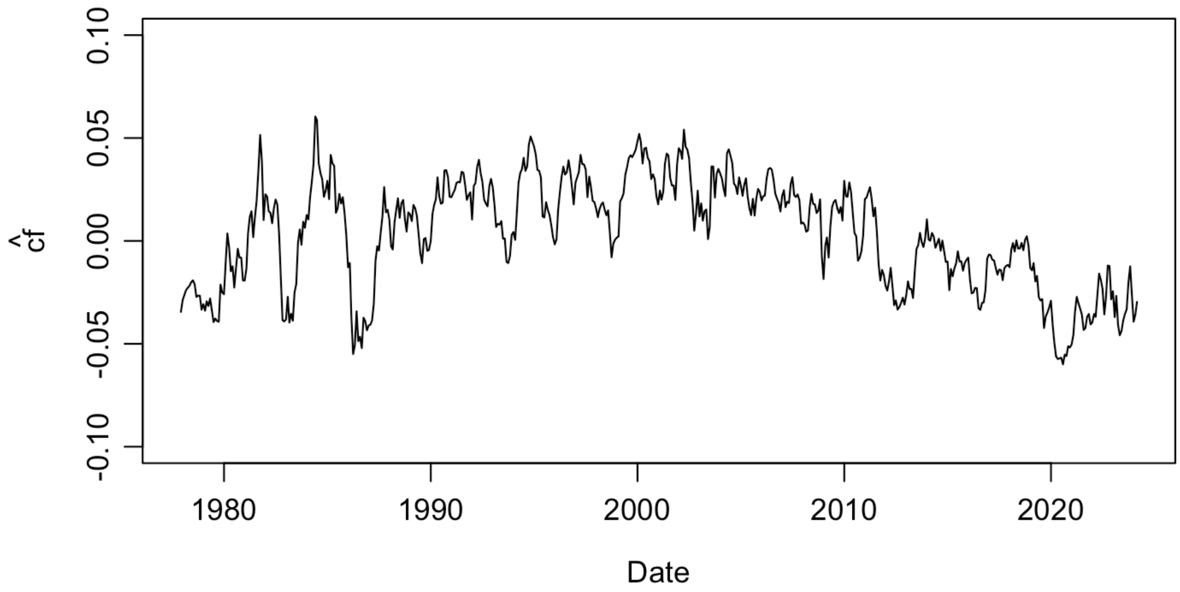

Figure 5.3: US return forecasting factor, $\hat{cf}$

The regression coefficients are given in Table 5.5. The β_1 row represents the estimated amounts by which bond returns will change per unit movement in the return forecasting factor across different horizons. For example, β_1 reaches 1.8367 for the 20-year horizon, signifying that the 20-year returns are very sensitive to changes in the return forecasting factor. Conversely, β_1 is only 0.1428 for the 1-year horizon, indicating that the 1-year returns are less responsive to fluctuations in the return forecasting factor. The β_1 values seem to increase as the time horizon increases, suggesting that the influence of the return forecasting factor on the bond returns increases over time.

Table 5.5: Regression coefficients for different horizons - factor model

	$rp^{(1)}$	$rp^{(2)}$	$rp^{(5)}$	$rp^{(7)}$	$rp^{(10)}$	$rp^{(20)}$
β_0	-0.0002	-0.0004	-0.0005	-0.0003	0.0004	0.0010
β_1	0.1428	0.2981	0.8428	1.2809	1.5987	1.8367

In-Sample adjusted R^2 Using OLS Regression Table 5.6 shows the adjusted R^2 values from Equation 4.13, which shows that up to 35.22% of the variation in future bond returns can be explained by the factor model.

Table 5.6: In-sample adjusted R^2 values across various maturities in US context - factor model

	$rp^{(1)}$	$rp^{(2)}$	$rp^{(5)}$	$rp^{(7)}$	$rp^{(10)}$	$rp^{(20)}$
adjusted R^2	0.1733	0.2303	0.3142	0.3413	0.3505	0.3522

Comparison of Full Model and Factor Model

Looking at Table 5.7, the OLS regression using the full model on in-sample data performs the best. However, the performance of the factor model is only slightly worse. This highlights the importance of the return forecasting factor, as it is a simpler yet high-performing model that uses a single factor to model returns across all maturities. The slightly lower adjusted R^2s are compensated by the factor model being much easier to interpret.

Table 5.7: Comparison of in-sample adjusted R^2 values in US context

adjusted R^2	$rp^{(1)}$	$rp^{(2)}$	$rp^{(5)}$	$rp^{(7)}$	$rp^{(10)}$	$rp^{(20)}$
OLS - full model	0.2555	0.2883	0.353	0.3833	0.3928	0.3900
Elastic Net - full model	0.2543	0.2868	0.3528	0.3528	0.3611	0.3823
OLS - factor model	0.1733	0.2303	0.3142	0.3413	0.3505	0.3522

5.2 SA CONTEXT

5.2.1 Cycles

Figure 5.4 shows the SA inflation data, trend inflation, the zero-coupon yields and the fitted values from the regression in Equation 4.4. The cycles, $c_t^{(n)}$, obtained from Equation 4.4 when using SA zero-coupon bond yields and SA trend inflation are then shown in Figure 5.5.

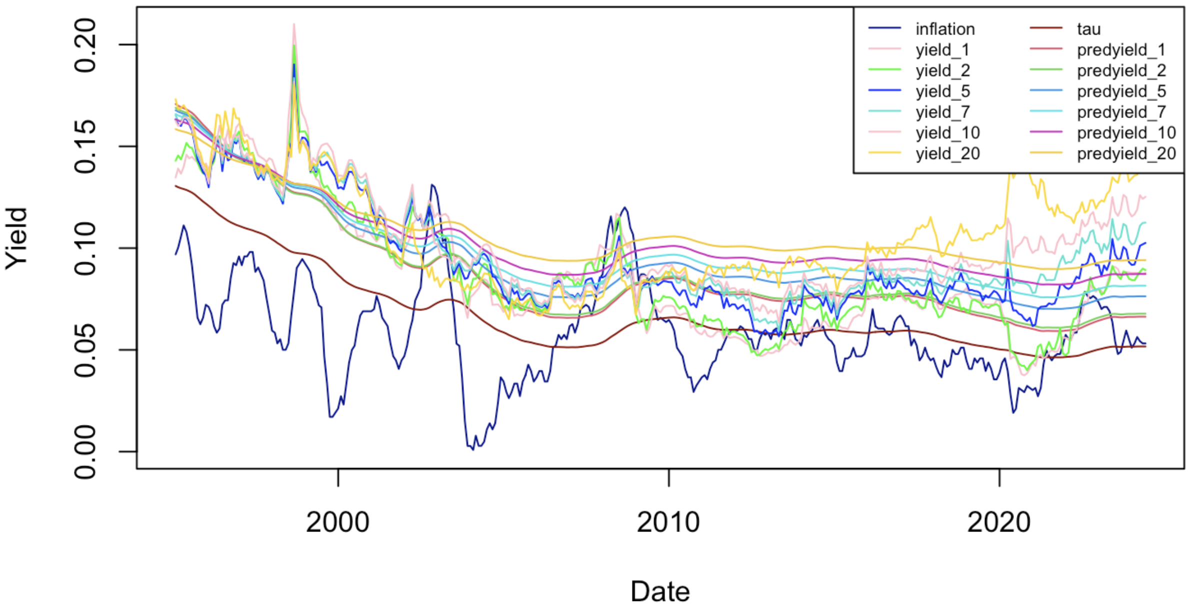

Figure 5.4: A time series of SA inflation, of the SA yields of maturities from 1- to 20-years, of the trend inflation (the variable τ_t^{proxy}) and the fitted yields from Equation 4.4

We see that cycles for the 1- and 2-year maturities seem to be mean-reverting, whereas the rest of the cycles do not seem to have this property. For example, the cycle for the 20-year maturity bond seems to trend upward from around 2005 onwards.

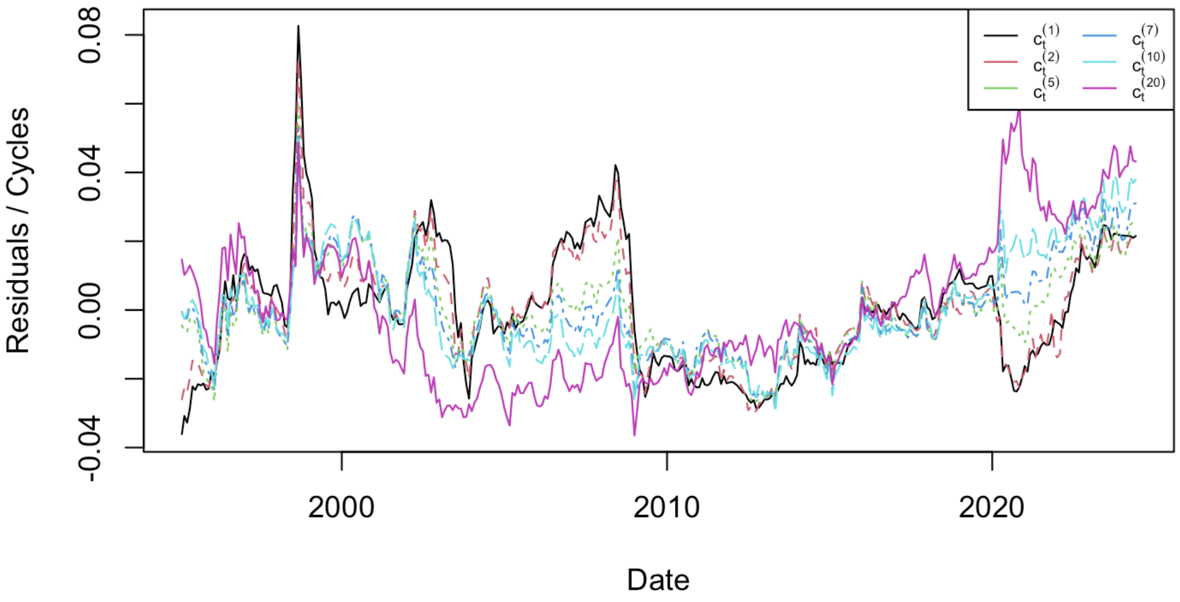

Figure 5.5: SA residuals/cycles, $c_t^{(n)}$ or $\varepsilon_t^{(n)}$, as in Equation 4.4

5.2.2 Cointegration

We would like to determine if a relationship exists between the SA zero-coupon yields and SA trend inflation, τ_t^{proxy}.

ADF Test for τ_t^{proxy} and Zero-Coupon Yields

Table 5.8 reports the p-values from the ADF test for τ_t^{proxy} and the zero-coupon yields. The ADF tests for τ_t^{proxy} and zero-coupon yields reveal p-values larger than 0.05; therefore, we cannot reject the null hypothesis and we conclude that both zero-coupon yields and τ_t^{proxy} have a unit root and are non-stationary.

Table 5.8: p-values from ADF test for τ_t^{proxy} and SA zero-coupon yields

No. of lags	τ_t^{proxy}	$y_t^{(1)}$	$y_t^{(2)}$	$y_t^{(5)}$	$y_t^{(7)}$	$y_t^{(10)}$	$y_t^{(20)}$
1	0.3509	0.3873	0.3890	0.5172	0.6068	0.7076	0.6388

ADF Test for Residuals

Table 5.9 shows that the p-values for the cointegration residuals (cycles) are smaller than 0.05 only up until term 2, therefore we reject the null hypothesis of a unit root up until term 2. We conclude that the cycles are stationary for term 1 and term 2 only, which suggests a stable, long-term relationship between the SA zero-coupon yields with maturity 1- and 2-years and τ_t^{proxy}. This also shows that the relationship between trend inflation and zero-coupon yields from 5- to 20-year maturities is not stable. The lack of cointegration between yields and trend inflation shows that the long-run Fisher effect does not necessarily hold in the South African market for longer maturity bonds. The lack of mean reversion among the cycles will likely decrease the predictive power of the models.

Table 5.9: p-values from ADF test for cointegrating residuals with SA data

No. of lags	$c_t^{(1)}$	$c_t^{(2)}$	$c_t^{(5)}$	$c_t^{(7)}$	$c_t^{(10)}$	$c_t^{(20)}$
1	0.0271	0.0215	0.0728	0.2226	0.4187	0.5219

5.2.3 Bond Returns and Predictability

Full Model

Again, recall the full model from Equation 4.10:

$$rp_{t+12}^{(n)} = \delta_0 + \sum_i \delta_i c_t^{(i)} + \varepsilon_{t+12}^{(n)}$$

Fitted values for each bond maturity were obtained using OLS and Elastic Net regression. The model's performance was evaluated using adjusted R^2 values.

In-Sample adjusted R^2 Using OLS and Elastic Net Regression Studying Table 5.10, we see that up to 32.57% of the variation in future bond returns in South Africa is explained by the full model.

Table 5.10: In-sample adjusted R^2 values for different maturities using OLS and Elastic Net regression in SA context - full model

adjusted R^2	$rp^{(1)}$	$rp^{(2)}$	$rp^{(5)}$	$rp^{(7)}$	$rp^{(10)}$	$rp^{(20)}$
OLS	0.3040	0.2838	0.3248	0.3241	0.3254	0.3195
Elastic Net	0.2731	0.2815	0.3228	0.3257	0.3250	0.3124

Regression Coefficients Again, we see the coefficients being shrunk to zero in Table 5.11 as a result of the Elastic Net regression.

Table 5.11: Comparison of OLS and Elastic Net regression coefficients for different horizons in the SA context - full model

	OLS						Elastic Net					
	$rp^{(1)}$	$rp^{(2)}$	$rp^{(5)}$	$rp^{(7)}$	$rp^{(10)}$	$rp^{(20)}$	$rp^{(1)}$	$rp^{(2)}$	$rp^{(5)}$	$rp^{(7)}$	$rp^{(10)}$	$rp^{(20)}$
δ_0	0.0024	0.0045	0.0087	0.0098	0.0105	0.0131	0.0023	0.0045	0.0087	0.0098	0.0104	0.0129
δ_1	0.7512	-0.3015	-1.0352	-1.8625	-3.2944	-5.1281	0.4369	-0.2531	-1.0356	-1.7334	-2.6066	-4.7115
δ_2	-0.6998	0.7486	-0.8502	0.1484	1.2612	3.2955	0.0000	0.7431	-0.8334	0.0000	0.0000	3.0681
δ_5	2.7799	3.9951	9.8049	7.6710	9.2321	10.2310	0.9227	3.6745	9.7086	7.1063	9.7268	5.8497
δ_7	-3.2679	-4.681	-5.3415	-1.1579	-4.8335	-8.3093	-0.8337	-4.7395	-5.2201	0.0000	-3.8682	0.0000
δ_{10}	0.8948	1.1064	-0.6542	-2.3203	0.9118	4.1964	0.0000	1.5908	-0.6943	-2.8961	0.0000	0.0000
δ_{20}	0.0852	0.2016	0.1078	-0.1040	-0.8006	-1.1367	0.0000	0.0000	0.1052	-0.0974	-0.7535	-0.9944

Factor Model

Recall the OLS regression from Equation 4.11:

$$\overline{rp}_{t+12} = \gamma_0 + \gamma_1 c_t^{(1)} + \gamma_2 \bar{c}_t + \varepsilon_{t+12}$$

An in-sample adjusted R^2 value of 0.1296 is obtained from the regression in Equation 4.11. This indicates that approximately 12.96% of the variation in the average SA returns across different maturities is explained by the variation in short rate expectations, $c_t^{(1)}$, and the variation in risk premia across all maturities, $\bar{c}_t$.

The fitted values from this regression correspond to the single return forecasting factor, $\widehat{cf}_t$, as shown in Equation 4.12:

$$\widehat{cf}_t = \hat{\gamma}_0 + \hat{\gamma}_1 c_t^{(1)} + \hat{\gamma}_2 \bar{c}_t$$

Which we regress in Equation 4.13:

$$rp_{t+12}^{(n)} = \beta_0^{(n)} + \beta_1^{(n)}\widehat{cf}_t + \varepsilon_{t+12}^{(n)}$$

Fitted Values & Regression Coefficients Figure 5.6 shows a plot of the return forecasting factor, $\hat{cf}$, over time. As mentioned before, a higher return forecasting factor indicates a greater likelihood of bonds rallying in the future, while a lower factor suggests an increased risk of bonds losing value (Cochrane and Piazzesi, 2005).

Figure 5.6: SA return forecasting factor, $\hat{cf}$

The regression coefficients are shown in Table 5.12. The regression coefficient, β_1, indicates the amount by which the traded bond will change per unit movement in the return forecasting factor. The regression coefficients increase as the time horizon increases, suggesting that the influence of the return forecasting factor increases as the bond horizon increases.

Table 5.12: Regression coefficients for different horizons - factor model

	$rp^{(1)}$	$rp^{(2)}$	$rp^{(5)}$	$rp^{(7)}$	$rp^{(10)}$	$rp^{(20)}$
β_0	-0.0008	-0.0009	0.0009	0.0008	0.0008	-0.0008
β_1	0.3750	0.6517	0.9834	1.1452	1.2045	1.6402

In-Sample adjusted R^2 Using OLS Regression Table 5.13 shows that at most 15.56% of the variation in future bond returns can be explained by the factor model.

Table 5.13: In-sample adjusted R^2 values across various maturities in SA context - factor model

	$rp^{(1)}$	$rp^{(2)}$	$rp^{(5)}$	$rp^{(7)}$	$rp^{(10)}$	$rp^{(20)}$
adjusted R^2	0.1304	0.1556	0.1246	0.1221	0.0961	0.1092

Comparison of Full Model and Factor Model

From Table 5.14, we observe that the full model performs significantly better than the factor model across all maturities. The lack of mean reversion amongst the cycles suggests the cycles do not necessarily tell us where yields are going to move in the future. Hence, compressing all the information across yields into one factor does not work well when the cycles are non-stationary.

Table 5.14: Comparison of in-sample adjusted R^2 values in SA context

adjusted R^2	$rp^{(1)}$	$rp^{(2)}$	$rp^{(5)}$	$rp^{(7)}$	$rp^{(10)}$	$rp^{(20)}$
OLS - full model	0.3040	0.2838	0.3248	0.3241	0.3254	0.3195
Elastic Net - full model	0.2731	0.2815	0.3228	0.3257	0.3250	0.3124
OLS - factor model	0.1304	0.1556	0.1246	0.1221	0.0961	0.1092

5.3 COMPARISON OF US AND SA RESULTS

Studying Table 5.15, it is observed that the full model still performs reasonably well in the SA context when compared to the US, however, the factor model in the SA context does not perform well at all. The 1- and 2-year SA cycles are stationary and therefore are still able to provide some predictive power in the full model as they are predictors on their own. In the factor model, however, the 2-year cycle is combined in an average with the rest of the higher maturity cycles. These higher maturity cycles are non-stationary and hence combining them takes away from the predictive power of the 2-year cycle. This is what diminishes the factor model's performance in the SA context.

Table 5.15: Comparison of in-sample adjusted R^2 values in SA and US contexts

adjusted R^2	$rp^{(1)}$	$rp^{(2)}$	$rp^{(5)}$	$rp^{(7)}$	$rp^{(10)}$	$rp^{(20)}$
US						
OLS - full model	0.2555	0.2883	0.353	0.3833	0.3928	0.3900
Elastic Net - full model	0.2543	0.2868	0.3528	0.3528	0.3611	0.3823
OLS - factor model	0.1733	0.2303	0.3142	0.3413	0.3505	0.3522
SA						
OLS - full model	0.3040	0.2838	0.3248	0.3241	0.3254	0.3195
Elastic Net - full model	0.2731	0.2815	0.3228	0.3257	0.3250	0.3124
OLS - factor model	0.1304	0.1556	0.1246	0.1221	0.0961	0.1092

This is further illustrated in Figure 5.7 and Figure 5.8. It is seen that in the full model, the SA adjusted R^2s are almost as high as the US adjusted R^2s, showing that allowing the stationary cycles to be predictors on their own still gives some predictive power to the model despite the non-stationary cycles being present. This effect is seen to diminish in the factor model when the cycles are combined into an average.

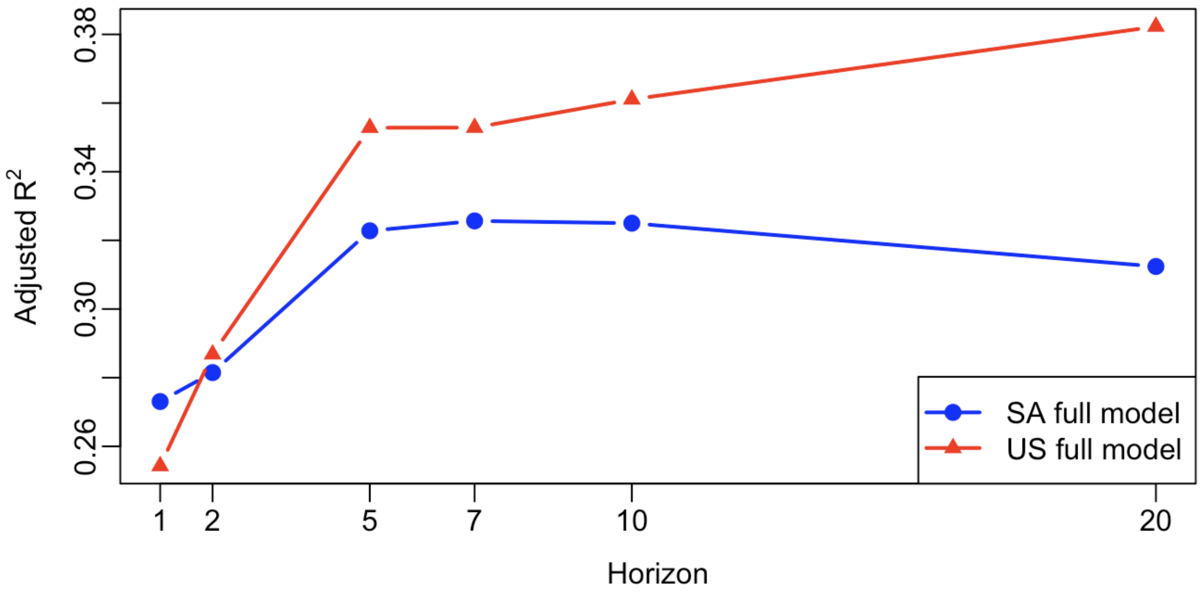

Figure 5.7: Comparison of adjusted R^2 values obtained from full model using Elastic Net regression in US and SA context

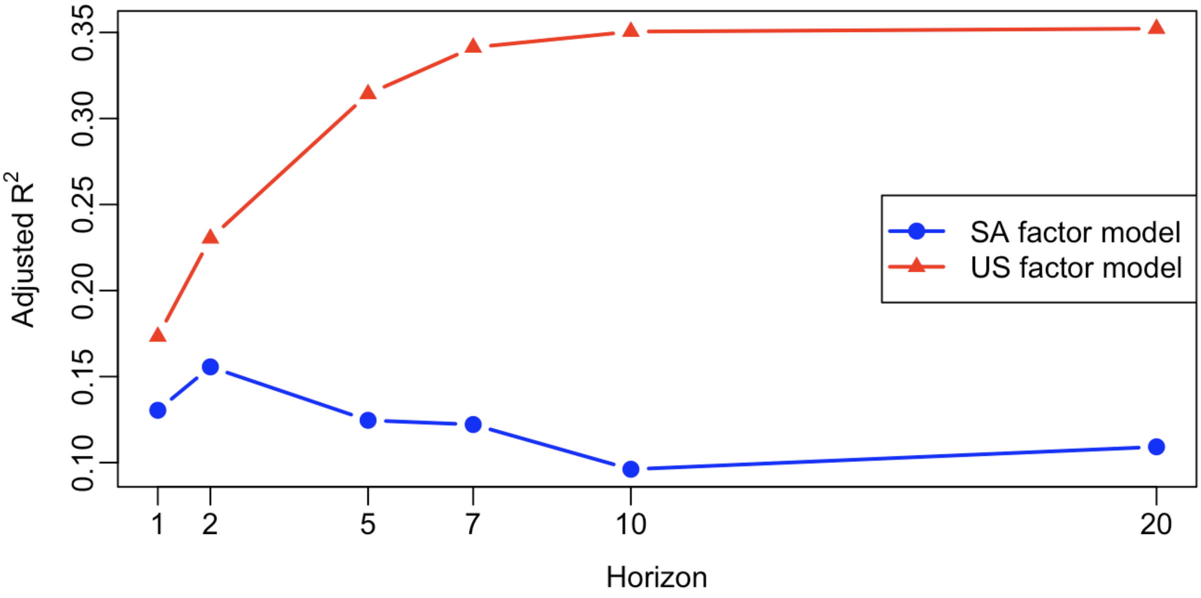

Figure 5.8: Comparison of adjusted R^2 values obtained from factor model in US and SA context

The results highlight that, while bond return predictability models such as those developed by Cieslak and Povala are useful, their effectiveness may vary significantly between different markets where cycles behave differently. It is shown that the models will work better in markets where the cycles exhibit mean-reverting behaviour.

This calls for further exploration of model adjustments or additional factors that could improve predictability in less developed bond markets like South Africa. This brings us to a combination of US and SA cycles in the following section.

5.4 COMBINING US AND SA CYCLES

We would like to improve the predictability of bond returns in South Africa as described in Section 4.3. Recall the regression equation:

$$rp_{t+12}^{(n)} = \delta_0 + \sum_i \delta_i c_t^{(i)} + \sum_i \delta_i^* c_t^{*(i)} + \varepsilon_{t+12}^{(n)}$$

where $c_t^{(i)}$ are the SA cycles and $c_t^{*(i)}$ are the US cycles.

Regression Coefficients Table 5.16 shows the regression coefficients obtained from the regression equation above. The US regression coefficients seem to be larger than the SA regression coefficients, indicating that the US zero-coupon cycles have a larger impact on the return of SA traded bonds than the SA zero-coupon cycles.

Table 5.16: Comparison of OLS and Elastic Net regression coefficients for different horizons

	OLS						Elastic Net					
	$rp^{(1)}$	$rp^{(2)}$	$rp^{(5)}$	$rp^{(7)}$	$rp^{(10)}$	$rp^{(20)}$	$rp^{(1)}$	$rp^{(2)}$	$rp^{(5)}$	$rp^{(7)}$	$rp^{(10)}$	$rp^{(20)}$
δ_0	0.0026	0.0049	0.0094	0.0104	0.0108	0.0122	0.0026	0.0049	0.0094	0.0103	0.0109	0.0125
δ_1	0.7434	-0.2288	-0.6691	-1.1375	-1.9521	-2.7834	0.7380	-0.2117	-0.6623	-1.1707	-1.4208	-1.5000
δ_2	-0.3326	1.1659	-0.7706	-0.2074	0.1785	1.0551	-0.3343	1.1306	-0.7770	0.0000	0.0000	0.0000
δ_5	2.2749	3.5263	9.9731	7.9418	9.4200	9.5867	2.2958	3.4142	9.9133	6.9971	5.5116	4.9605
δ_7	-2.2594	-3.7620	-5.6193	-1.4216	-4.9069	-5.7437	-2.2820	-3.3260	-5.5050	0.0000	0.0000	0.0000
δ_{10}	-0.0049	0.3969	-0.1160	-1.9767	0.9987	0.7390	0.0045	0.0000	-0.2068	-2.7218	0.0000	0.0000
δ_{20}	0.4912	0.6454	0.4153	0.6362	0.4684	2.4822	0.4932	0.7400	0.4427	0.7117	0.0114	1.6592
δ_1^*	-0.6609	-1.5059	-4.1174	-5.5058	-6.8886	-6.9250	-0.6074	-1.4058	-3.9088	-5.3288	-6.5853	-6.5217
δ_2^*	2.5589	4.5466	9.9205	12.2068	14.5505	10.9473	2.4328	4.2487	9.5069	11.9042	14.2882	11.6074
δ_5^*	-7.2145	-11.8868	-19.4710	-20.6812	-23.4626	-10.1355	-6.9583	-11.2245	-19.0757	-21.0434	-25.1922	-19.4478
δ_7^*	4.2489	4.8732	1.4890	-1.5358	-2.3206	-11.0532	4.0083	4.7251	1.6261	0.0000	0.0000	0.0000
δ_{10}^*	-0.0488	4.4487	17.5735	21.7732	25.9086	19.3929	0.0000	3.7773	17.0784	20.3252	25.9783	19.2136
δ_{20}^*	1.2091	-0.3191	-4.8029	-4.8770	-5.3188	2.7507	1.2197	0.0639	-4.6225	-4.4872	-6.0336	0.1091

In-Sample adjusted R^2 Using OLS and Elastic Net Regression From Table 5.17, we see that 46.77% to 57.98% of the variation in the SA returns can be explained by the SA and US cycles. The in-sample adjusted R^2 values obtained in this regression are significantly higher than the in-sample adjusted R^2 values obtained in previous regressions using SA data alone.

Table 5.17: In-sample adjusted R^2 values obtained from combining the US and SA cycles using OLS and Elastic Net regression

adjusted R^2	$rp^{(1)}$	$rp^{(2)}$	$rp^{(5)}$	$rp^{(7)}$	$rp^{(10)}$	$rp^{(20)}$
OLS	0.5034	0.4681	0.4828	0.5097	0.5370	0.5798
Elastic Net	0.5034	0.4677	0.4812	0.5119	0.5329	0.5727

These results show that SA bond returns are driven by local factors as well as economic factors in the US and that combining both US and SA cycles improves the predictability of SA bond returns.

5.5 OUT-OF-SAMPLE RESULTS

Out-Of-Sample R^2

Out-of-sample R^2 measures how well a model performs on unseen data and its values lie on the interval $(-\infty, 1]$, instead of $[0, 1]$ for in-sample R^2 (Hawinkel *et al.*, 2024).

In both the US and SA contexts, we found that all models produced a negative out-of-sample R^2 as seen in Table 5.18. These results suggest that despite the high variation captured by these models in-sample, they are not robust when it comes to accurately forecasting the future returns.

Table 5.18: Comparison of out-of-sample R^2 values in US, SA and SA with SA & US cycles

R^2	$rp^{(1)}$	$rp^{(2)}$	$rp^{(5)}$	$rp^{(7)}$	$rp^{(10)}$	$rp^{(20)}$
US						
OLS - full model	-0.2219	-0.2226	-0.1922	-0.1373	-0.1272	-0.1668
Elastic Net - full model	-0.1582	-0.1537	-0.1423	-0.0971	-0.1116	-0.1299
OLS - factor model	-0.1647	-0.1510	-0.1310	-0.0948	-0.0945	-0.0807
SA						
OLS - full model	-0.6066	-0.4438	-0.5451	-0.582	-0.7594	-1.1556
Elastic Net - full model	-0.7523	-0.4079	-0.5298	-0.6109	-0.684	-1.1591
OLS - factor model	-0.5582	-0.5564	-0.8502	-0.9428	-1.1586	-1.4772
SA with SA & US Cycles						
OLS	-0.4230	-0.3088	-0.5485	-0.4802	-0.5163	-0.7979
Elastic Net	-0.3700	-0.2765	-0.5891	-0.4128	-0.4913	-0.7848

The substandard forecasting performance of the models brings to question the usefulness of them in an investment decision making setting. One might think that this poor performance would imply that the models are redundant in real-world applications. This is not necessarily the case.

In an investment setting, we are not always concerned with what the exact future return is going to be. We are also concerned with whether the returns are going to be positive or negative. In other words, we are interested in which direction the market is heading. Hence, we also look at whether or not the models presented can be useful in predicting the market movement. To evaluate the accuracy of the model in predicting market movements, we look at what proportion of the time they correctly predict the direction of the return.

Accuracy

Studying Table 5.19, the accuracy in the US context ranges from 60.24% to 66.67%, suggesting that the models in the US context would correctly predict the direction of returns 60.24% to 66.67% of the time. The accuracy for each model at each maturity is very similar.

Moreover, the accuracy in the SA context ranges from 43.06% to 55.56%, suggesting that the models in the SA context would correctly predict the direction of returns 43.06% to 55.56% of the time. It seems that there is no significant difference between the accuracy between the full models, using different regressions, and the factor model. However, the accuracy is substantially lower in the SA context compared to the US context.

Lastly, when combining SA and US cycles to predict SA Treasury bond returns, the accuracy ranges from 57.41% to 64.35%, which is a great improvement from just using the SA cycles. The accuracy results tell us that, despite the negative out-of-sample R^2s, the models presented can still be a useful signal to make investment decisions. To illustrate this further, we look at a 7-year bond in both the US and SA context.

Table 5.19: Comparison of accuracy of models in US, SA and SA with SA & US cycles

	$rp^{(1)}$	$rp^{(2)}$	$rp^{(5)}$	$rp^{(7)}$	$rp^{(10)}$	$rp^{(20)}$
US						
OLS - full model	0.6214	0.6667	0.6476	0.6500	0.6452	0.6524
Elastic Net - full model	0.6048	0.6619	0.6452	0.6643	0.6571	0.6595
OLS - factor model	0.6024	0.6643	0.6405	0.6405	0.6429	0.6643
SA						
OLS - full model	0.5093	0.4861	0.5046	0.4861	0.4861	0.4861
Elastic Net - full model	0.4491	0.4306	0.5556	0.5000	0.5046	0.4861
OLS - factor model	0.4537	0.5139	0.4769	0.4861	0.4676	0.4861
SA with SA & US Cycles						
OLS	0.6343	0.5880	0.6157	0.6435	0.6019	0.5741
Elastic Net	0.5880	0.5833	0.6435	0.6250	0.6296	0.5926

Evaluation of 7-year Treasury Bond in US and SA Context

Figure 5.9 and Figure 5.10 show the predicted versus observed returns using Elastic Net regression. To illustrate how well the model predicts market movements, we look at some more performance metrics.

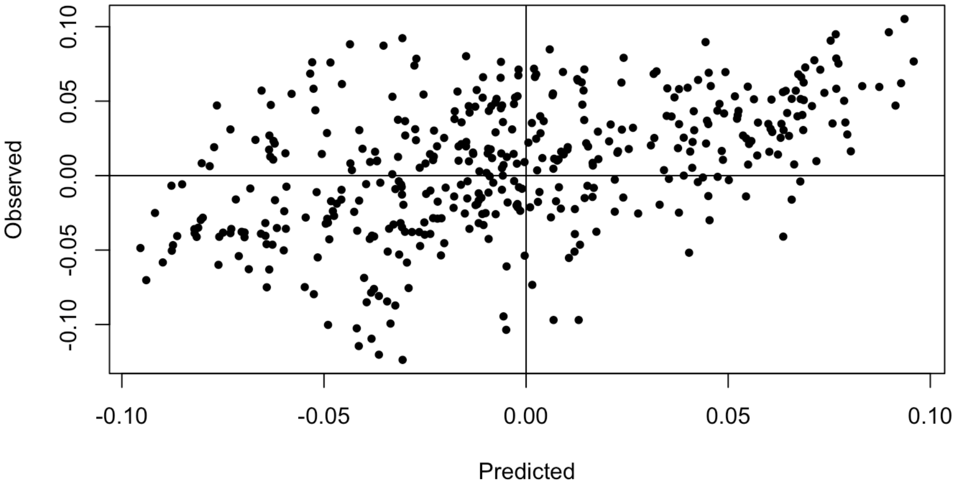

Figure 5.9: US 7-year Treasury bond returns: predicted vs. observed using Elastic Net regression - full model

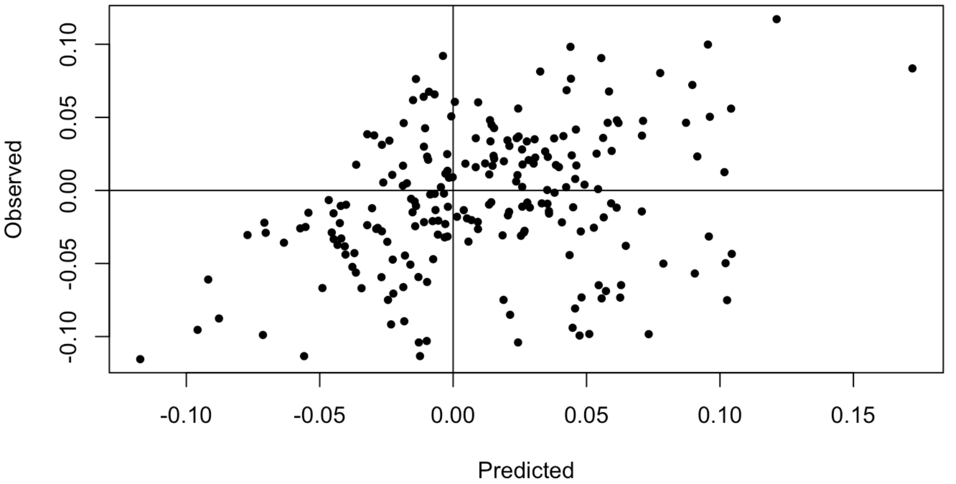

Figure 5.10: SA 7-year Treasury bond returns: predicted vs. observed using Elastic Net regression and combined US & SA cycles

Precision is the probability that a return is positive given it was predicted positive. Recall is the probability of predicting a positive return given it is positive. The F1 score is the harmonic mean of the precision and recall values. Since the harmonic mean gives more weight to low values, the F1 score for the model will only be high if both the precision and recall are high (Géron, 2019).

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}$$

$$\text{F1 Score} = \frac{2 \times \text{precision} \times \text{recall}}{\text{precision} + \text{recall}}$$

where TP, FP, TN and FN are as defined in Section 4.4.

Table 5.20: Performance metrics for 7-year Treasury bond using Elastic Net regression: US vs. SA with US & SA cycles

	US	SA
Accuracy	0.6643	0.6250
Precision	0.7841	0.5620
Recall	0.5726	0.7083
F1	0.6619	0.6267

A low precision value means the model often incorrectly predicts a positive return, which can lead to losses. This is represented by the observations in the bottom right quadrant of Figure 5.9 and Figure 5.10. In the US context, the precision value shows that the model predicted a return would be positive and it was actually positive 78.41% of the time, whereas the combined model in the SA context predicted a positive return and the observed return was only positive 56.20% of the time. They are both correct over 50% of the time; however, the model performs better by this metric in the US context. Moreover, a low recall value means the model often predicts negative returns when the returns are actually positive. This represents an opportunity cost, illustrated by the top left quadrants of Figure 5.9 and Figure 5.10. We see that the recall was much better in the SA context. The difference in precision and recall values for the different markets shows that in the US context, the model is more likely to miss an investment opportunity and in the SA context, the model is more likely to incorrectly predict a positive return. Despite this, these models can still be a useful tool to make investment decisions in both markets, as shown by the accuracy and F1 score.

CHAPTER 6

CONCLUSION

This research has investigated the predictability and explainability of Treasury bond returns, building on the model explained by Cieslak and Povala (2011). By applying this methodology to the US market, we have found that the full model as well as the factor model explain a large proportion of the variance of returns on bonds that are actively traded in the market. The full model performs almost as well in the South African context; however, the factor model does not perform well at all. This results from not all the cycles being stationary, with a long-term relationship between yields and trend inflation not present at all maturities.

To better model the South African bond returns, we proposed a new model based on the idea that the US economy has a strong effect on emerging markets like South Africa. Combining US and South African cycles in the regression analysis demonstrates a significant improvement, with adjusted R^2 values ranging from 0.4677 to 0.5798. This suggests that global economic factors have a significant impact on local markets, bringing attention to the connection between bond markets around the world.

After making predictions on out-of-sample data, the out-of-sample R^2s were negative in both the US and South African contexts, showing that all the models used did not exhibit good forecasting abilities and are hence not robust out-of-sample.

Despite this, the models were still shown to be useful in making investment decisions, with the full model in the US context and the combined model in the SA context providing an accurate signal of which direction the market will move in the future at least 60% of the time.

Overall, these insights have extended upon the work done by Cieslak and Povala (2011) and have shown that the effectiveness of these models varies significantly between the developed US market and the emerging South African market. The performance of the model combining the US and South African cycles to model South African Treasury bond returns resulted in a massive improvement over using the full model and factor model in South Africa. Despite the poor forecasts from all the models, they were still shown to be useful tools in an investment setting.

Further research on this topic could investigate how these models perform on other emerging markets, providing even more insight into the models' generalisability in different economic contexts. Moreover, looking at different ways of modelling trend inflation could improve the performance of the models. Another important area of further research is evaluating the effect of the sample size on model performance, as well as looking at different time frames of returns. Lastly, more advanced machine learning techniques could be looked at to potentially improve the out-of-sample forecasts.

REFERENCES

Adrian, C. and Moench (2013). Pricing the term structure with linear regressions. *Journal of Financial Economics*, vol. 110, no. 1, pp. 110–138.

Brooks, J. (2021 July). What drives bond yields? `https://images.aqr.com/-/media/AQR/Documents/Whitepapers/Drivers-of-Bond-Yields.pdf`. Accessed: 2024-08-13.

Cieslak and Povala (2011). Understanding bond risk premia. *Unpublished working paper. Kellogg School of Management, Evanston, IL*, pp. 677–691.

Cochrane and Piazzesi (2005). Bond risk premia. *American economic review*, vol. 95, no. 1, pp. 138–160.

Dias, D.A. and Marques, C.R. (2010). Using mean reversion as a measure of persistence. *Economic Modelling*, vol. 27, no. 1, pp. 262–273.

Fedderke, J.W. (2021). The south african–united states sovereign bond spread and its association with macroeconomic fundamentals. *South African Journal of Economics*, vol. 89, no. 4, pp. 499–525.

Federal Reserve Bank of St. Louis (2024). Consumer price index for all urban consumers: All items less food and energy in u.s. city average. Accessed: 2024-09-08.
Available at: `https://fred.stlouisfed.org/series/CPILFESL`

Fisher, I. (1907). *The Rate of Interest: Its nature, determination and relation to economic phenomena.* Macmillan.

Friedman, J., Hastie, T. and Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*, vol. 33, no. 1, p. 1.

Géron, A. (2019). Hands-on machine learning with scikit-learn, keras & tensorflow farnham. *Canada: O'Reilly.*

Hawinkel, S., Waegeman, W. and Maere, S. (2024). Out-of-sample r 2: Estimation and inference. *The American Statistician*, vol. 78, no. 1, pp. 15–25.

James, G., Witten, D., Hastie, T., Tibshirani, R. *et al.* (2013). *An introduction to statistical learning*, vol. 112. Springer.

Romeo, G. (2020). Chapter 13—data analysis for business and economics. *Elements of Numerical Mathematical Economics with Excel; Academic Press: Cambridge, MA, USA*, pp. 695–761.

Stigler, M. (2020). Nonlinear time series in r: Threshold cointegration with tsdyn. In: *Handbook of Statistics*, vol. 42, pp. 229–264. Elsevier.

Sztepanacz, J.L. and Houle, D. (2024 February). Regularized regression can improve estimates of multivariate selection in the face of multicollinearity and limited data. *Evolution Letters*, vol. 8, no. 3, pp. 361–373. ISSN 2056-3744. `https://academic.oup.com/evlett/article-pdf/8/3/361/58955441/qrad064\_suppl\_supplementary\_material.pdf`.
Available at: `https://doi.org/10.1093/evlett/qrad064`

Taylor, J.B. (1993). Discretion versus policy rules in practice. In: *Carnegie-Rochester conference series on public policy*, vol. 39, pp. 195–214. Elsevier.