Handelshøyskolen BI

GRA 19703 Master Thesis

Thesis Master of Science 100% - W

Predefinert informasjon

Startdato:	09-01-2023 09:00 CET	**Termin:**	202310
Sluttdato:	03-07-2023 12:00 CEST	**Vurderingsform:**	Norsk 6-trinns skala (A-F)
Eksamensform:	T		
Flowkode:	202310||11184||IN00||W||T		
Intern sensor:	(Anonymisert)		

Deltaker

Navn:	Andre Nikolai Nilsen og Dmitrij Minasow

Informasjon fra deltaker

Tittel *:	The Influence of Yield Curve Factors on the Conditional Distribution of Equity Risk Premia
Navn på veileder *:	Paolo Giordani

Inneholder besvarelsen konfidensielt materiale?:	Nei	**Kan besvarelsen offentliggjøres?:**	Ja

Gruppe

Gruppenavn:	(Anonymisert)
Gruppenummer:	195
Andre medlemmer i gruppen:	

GRA 19703

The Influence of Yield Curve Factors on the Conditional Distribution of Equity Risk Premia

---

Master Thesis

by
Andre Nikolai Nilsen & Dmitrij Minasow
MSc in Quantitative Finance

Supervised by
Paolo Giordani

Abstract

This paper examines the efficacy of how certain yield curve factors, namely the level, slope, and curvature, can predict the shape of cumulative return distributions in equity markets. We utilise the GAMLSS framework to analyse distributional characteristics and examine the impact yield curve factors have on equity returns. Our findings indicate that the slope of the yield curve is the most influential factor affecting the shape of the return distribution, compared to other factors such as level and curvature. Specifically, as the slope becomes increasingly upward-sloping, the return distribution approaches symmetry, while a lower slope leads to the distribution becoming more negatively skewed. The results highlight the importance of accurate distributional assumptions in estimating risk metrics for making informed investment decisions. This is demonstrated by comparing several models based on different distribution families and incorporating more complexity to capture the characteristics in financial time series. Additionally, the predictive power of the slope on the shape of the distribution diminishes in US markets after the mid-1980s, consistent with existing academic literature that suggests a diminished effect in the slope's ability to forecast output growth after the mid-1980s. However, we report evidence suggesting that the slope remains relevant for predicting the shape of the distribution in other developed markets.

Keywords*: Equity risk premia, yield curve factors, GAMLSS, distributional modelling, risk management, skewness.*

Oslo, June 28, 2023

Acknowledgement

We sincerely thank our supervisor, Paolo Giordani, for his consistent guidance and support while writing this master's thesis.

Table of Contents

Table of Equations

1. Introduction

Extensive research has been devoted to the prediction of equity returns, with researchers such as (Miller & Modigliani, 1961), (Fama & French, 1993) and (Campbell & Shiller, 1998) making significant contributions in this field. Through their work, they have defined variables that they claim to be persistent for predicting expected equity risk premia. More recently, in his book, Antti Ilmanen (2011) combines findings from over 20 years of research on this topic. Despite being touted in academic literature as having predictive power for the equity premium, most variables have performed poorly both in-sample and out-of-sample (Welch & Goyal, 2008). Nonetheless, many variables have an asymmetric effect on the return distribution, affecting lower, central, and upper quantiles differently (Cenesizoglu & Timmermann, 2008).

In this thesis, we aim to investigate the efficacy of the yield curve in predicting conditional excess return distributions in the equity market. While most literature in the field focuses on forecasting expected rates of return, the distribution of returns, and its conditional and time-varying nature, remains a largely undiscussed topic.

The study of return distributions is essential for many financial professionals. Those include portfolio managers, who can leverage this information to make informed investment decisions; risk managers, who can evaluate the risk inherent in different investments; financial advisors, who can help clients grasp the possible outcomes of their investments; and regulators, who can gauge the effects of financial instruments and products on the financial system at large.

An important question that needs to be addressed is the ability of the yield curve to forecast the mean and, if not, whether there is a need to delve beyond the mean to comprehend the changes in the distribution's shape. It is also essential to determine if these predictors provide meaningful information on risk measures during periods of relatively high or low yield curve factor readings. We refer to the shape as the distribution's asymmetry, which we use as a proxy for skewness.

Given the ongoing changes in the global interest rate environment, we find this topic particularly relevant. Since the mid-1980s, markets have been experiencing a long-term cycle of decreasing interest rates, with some countries like Japan, Switzerland, and the Eurozone experiencing negative interest rates. The future of interest rates remains uncertain, and we don't know if we will return to a period of long-term increasing rates as we did before the mid-1980s. Regardless, by conducting this study, we can better prepare for any scenario or, at the very least, deepen our understanding of the equity return distribution conditional on changes in the yield curve.

We utilise Generalized Additive Models for Location, Scale, and Shape (GAMLSS) to address the time-varying nature and the asymmetric effects on the return distributions. This framework provides flexibility in choosing distributional assumptions, which is crucial when working with fat-tailed and asymmetric data. Furthermore, it enables us to model the distribution of the equity risk premia based on the state of the yield curve factors, facilitating the comparison of different interest rate environments and their impact on the return distribution.

Our approach begins with an Ordinary Least Squares regression (OLS), followed by a discussion of its limitations. We then transition to a GAMLSS model with a normal distribution assumption, demonstrating that this model produces results identical to the OLS. To account for fatter-than-normal tails, we extend our model by introducing a student-t distribution, which still retains symmetry. Furthermore, we introduce a Skewed-t distribution developed by Fernandez and Steel (1998) to capture asymmetry. Based on this distribution, we construct a Constant-Skew-T model that captures asymmetry by modelling it as a constant but freely estimated. Lastly, we implement the Conditional-Skew-T model to model the shape parameter as a function of yield curve factors while keeping the other distribution parameters constant.

This thesis aims to contribute to the existing literature, particularly within the framework established by Giordani and Halling (2019). By moving beyond the mean of equity returns, we aim to provide a more nuanced understanding of how the conditional distribution of equity returns fluctuates in response to changes in yield curve factors.

One of the key questions we seek to answer is whether fluctuations in yield curve factors during periods of high or low readings lead to significant changes in risk metrics. Moreover, we aim to challenge the convention of using the normal assumption in dealing with financial return time series, arguing that this assumption is often violated. We illustrate the impact of miss-specifying distribution assumptions by showing how much this influences risk metrics such as standard deviation, Value-at-Risk, and Expected Shortfall.

Finally, we aim to examine whether the yield curve factors as predictors for the return distribution have undergone any notable changes when comparing a sample before and after the mid-1980s. Through these objectives, our study aims to broaden the understanding of the dynamics between the equity risk premia and changing monetary policy regimes.

We structure our thesis in the following way: Chapter 2 starts our discussion by exploring the related literature to our research question. Chapter 3 details the data we are using, explaining the construction of yield curve factors, presenting descriptive statistics, and showcasing the preliminary analysis along with the limitations of using the OLS regression for our research question. Chapter 4 carefully walks through the methodology, introducing the GAMLSS model, model specifications, a comparison of models, interpretation of results, and computation of risk metrics. Our primary findings are detailed in Chapter 5 and concentrate on our parameter estimates, predictive distributions, their implications, a robustness check using international evidence, and lastly, highlighting the limitations of our study. The thesis concludes with Chapter 6, which presents the conclusion of our study. Chapter 7 provides a list of references, and finally, Chapter 8 serves as the appendix, showing figures that were not included in the main body of the paper.

2. Literature Review

2.1: Interest Rate Indicators and Equity Market Returns

The shape of the yield curve is closely related to business cycles, credit cycles, and monetary policy cycles. When the spread between the US 10-year and the US 3-month inverts, it has frequently foreshadowed economic recessions, but during such downturns, the curve tends to rapidly become steeper and peak near business cycle troughs (Ilmanen, 2011). The countercyclical property of yield curves could shed light on why it can be effective in anticipating short-term stock returns: when the business cycle reaches its low point, high premia demand results in a steep yield curve, but when the business cycle reaches its peak, low premia requirements lead to an inverted yield curve. Similarly, Chen (1991) finds a significant relationship between the term premium, short-term rate, and future economic growth. Most importantly, he finds that the expected economic activity, in turn, forecasts the market excess returns.

In their paper, "Business Conditions and Expected Returns on Stocks and Bonds", Fama and French (1989) find empirical evidence that the term spread is closely related to shorter-term business cycles. In particular, the term spread, and the component of expected returns it tracks, are low around measured business-cycle peaks and high near troughs. Fama and French (1989) find that the relationship is strongest for 1-year return lags and tends to decay thereafter. They suggest that the spread tracks a term or maturity premium in expected returns that are similar for all long-term assets and that the premium compensates for exposure to *discount-rate shocks* that affect all long-term securities. Contemporary research also confirms Fama and French's findings that the term structure of interest rates predicts excess stock returns, as well as excess returns on T-bills and bonds (Campbell, 1987), (Schwert, 1990).

More recent research also finds a positive relationship between the term spreads and future economic output. Ang et al. (2006) and Estrella (2005) show that the slope of the yield curve forecasts output growth and inflation. Interestingly, several studies find that the spread has forecast output growth less accurately since the mid-1980s, which some attribute to greater stability of output growth and other key macroeconomic data (D'Agostino et al., 2006), (Dotsey, 1998), and (Estrella et al., 2003). Other researchers also find empirical evidence of the term spread's predictive power outside the US. For example, Plosser and Rouwenhorst (1994) find that term spreads are useful for predicting GDP growth in Canada and Germany, as well as the United States, but not in France or the United Kingdom.

Supporting Fama and French's argument that the discount-rate shock premium affects all long-term assets, Cochrane and Piazzesi (2005) provide a compelling argument that stocks can be viewed as a long-term bond plus cash-flow risk. Therefore, any variable that forecasts bond returns should also predict stock returns. Following this logic, Cochrane and Piazzesi construct a single-factor model that incorporates the yield curve's level, slope, and curvature factors, which are commonly used in bond prediction (Veronesi, 2010). They find the single-factor model useful also for stock prediction. However, Cochrane and Piazzesi's single factor is estimated on the whole sample and is unknown ex-ante. Therefore, we construct the level, slope and curvature and test their impact on the stock returns separately.

Bernanke and Gertler (1995) accredit the predictive power of monetary indicators to the *credit channel* of monetary policy transmission. A tighter monetary policy leads to a reduced and costlier bank loan supply that, in turn, impacts future cash flows and stock returns. Supporting the credit channel theory, (Chava & Purnanandam, 2011) argues that banks' inability or unwillingness to extend credit to firms causes firms to reduce investment in some positive NPV investment projects leading to a drop in firm value. Further, (Chava et al., 2015) present evidence that tight lending standards predict future returns through cash flows, and the market does not immediately impound this information. This justifies our selection of the term structure level, used as an estimate for financing costs, as one of the variables for prediction.

2.2: Non-normality and Modelling Asymmetry

While a substantial amount of research exists documenting the effect of interest rate indicators on the mean of expected returns, classic research conducted by Fama, French and Campbell fails to recognise the time-varying impact of these variables on the return distribution's shape. More importantly, the framework under which the factors have been studied assumes a normal distribution, implying a constant and symmetric distribution. However, studies such as (Fama, 1965), (Mandelbrot, 1963) and (Mandelbrot & Taylor, 1967), among others, have provided evidence that the distribution of equity returns has fatter tails and does not conform to a normal distribution. Furthermore, not only do equity returns exhibit fatter tails, but aggregate stock market returns also display negative skewness (Albuquerque, 2012).

The explanation for the time-varying risk premia and non-normality of equity returns is multifaceted, with several theories offering insights. Non-normality, represented by longer-than-normal tails, can arise from periods of volatile equity returns, whilst return asymmetry is observable through rapid price declines during stock market crashes. Research indicates that the equity risk premia (ERP) fluctuate throughout the business cycle, and the ERP refers to the (expected or realised) return of a broad equity index in excess over some non-equity alternative (Ilmanen, 2011). The ERP can also be defined as the reward for bearing losses during bad times, defined by low consumption growth, disasters, or long-run risks (Ang, 2014). Ilmanen (2011) argues that there is a strong connection between the time-varying growth premium and the time-varying equity risk premium.

Ilmanen (2011) further argues that aggregate liquidity measures predict future market returns, which is aligned with the concept of time-varying liquidity premia. The presence of excess skewness, caused by sudden decreases in stock prices, can be explained through the liquidity spiral concept proposed by Brunnermeier & Pedersen (2009). This concept states that decreased liquidity, falling stock prices, margin calls, and forced selling can create a vicious cycle that leads to further declines in stock prices. Additionally, the phenomenon of "flight to safety" suggests that investors' risk aversion changes over time, causing abrupt drops in the prices of risky assets (Yoon, 2015). Lastly, Ilmanen (2011) indicates that a high level of illiquidity tends to occur during bear market conditions and mildly predicts high medium-term future returns. However, this effect seems to be stronger for small-cap firms.

Giordani & Halling (2019) document the link between valuation levels and the shape of the distribution of cumulative total returns. Specifically, when valuations levels are high, they find that the return distribution becomes more asymmetric and negatively skewed. Contrastingly, when valuations levels are low, they document a roughly symmetric return distribution. While Giordani and Halling's study documents the impact of valuation ratios on return distributions, we extend their modelling framework to explore the relationship between the yield curve factors and equity return distributions. To accomplish this, we utilise the GAMLSS (Generalized Additive Models for Location, Scale, and Shape) framework, which offers a convenient and novel way to examine the time-varying effect on distribution parameters.

3. Data Description

3.1 Data Sources

This section provides an overview of the data sources used to construct the slope, level, and curvature factors of the yield curve, data for equity market returns and any additional control variable incorporated in the thesis. All bond yields and returns are presented on a daily basis.

The daily data for the 10-year, 5-year and 3-month market yields on US Treasuries were retrieved from Bloomberg. The daily yields for the 10-year and 5-year Government Bond commences on the 2nd of January 1962, while the 3-month market yield for Treasury Bills starts on the 1st of January 1934. For equity market excess returns, data is available since 1926. These returns encompass the valueweighted performance of all firms listed on the New York Stock Exchange (NYSE), the American Stock Exchange (AMEX), and the NASDAQ. The daily market excess returns data were gathered from the Fama-French database. The data analysis for our main results starts from the 2nd of January 1962 to match the earliest available market yield.

Lastly, to ensure the robustness of our findings, we gather data for international markets. From Bloomberg, we retrieve daily data for the 10-year yield, the 5-year yield, the 2-year yield, and the total return indices for the UK, France, Germany, and Australia. This data was synchronised to the 29th of May 1992 to achieve a comparable sample size for all international markets.

3.2 Definitions and Variable Construction

In this section, we present the response variables, cumulative 12- and 6-month log excess returns, and our covariates: namely, the level, slope, and curvature. The holding period of 12 months is chosen because of its documented strong signal (Campbell, 1987), (Schwert, 1990), while six months is chosen to capture short-term business cycles described by (Fama & French, 1989).

We assemble the traditional yield curve using three distinct market yields. To maximise our sample size, we select the 3-month yield as the short-term rate, the 5-year yield as a medium-term rate, and the 10-year yield as the long-term rate. A common example of yield curve factor construction can be found in Veronesi (2010). While Veronesi employs the 1-month yield as the short-term rate, our model utilises a 3-month yield, a measure more commonly used in yield curve construction across academic literature.

Level

The level of the yield curve represents the average interest rate across all maturities, providing a general rate of return that investors can expect from a bond, irrespective of its term. Thus, the level represents the general cost of financing for firms in the economy as a proxy for the credit channel described by Bernanke and Gertler (1995).

$$LVL = \ln(\frac{1}{3}(y_{10yr} + y_{5yr} + y_{3M})) \tag{3.1}$$

Slope

The slope of the yield curve measures the spread between short-term and long-term interest rates. A high (low) slope indicates a high (low) term premium on a 10-year Treasury Bond compared to the 3-month Treasury Bill.

$$SLP = \ln(y_{10y} - y_{3M}) \quad 3.2$$

Curvature

Lastly, the yield curve's curvature represents the difference between mid-term and the average of short-term and long-term yields (constructed as a butterfly spread). High (low) curvature values indicate that the mid-term yield is relatively high (low) compared to the 10-year yield and 3-month yield. High values represent a concave term structure, while low curvature values represent a convex term structure.

$$CRV = \ln(-y_{10yr} + 2 * y_{5yr} - y_{3M}) \quad 3.3$$

3.3 Data Cleaning and Transformation

The data cleaning and transformation process was conducted in a two-step manner. First, data was gathered from various sources and consolidated in Excel, then imported to R for additional processing. Within the R environment, the market excess return data was synchronised by trading dates with the data from US Treasury yields. Any missing daily yield observations were addressed using a forward-filling technique, which involves carrying the last observation forward.

All variables were converted to logs, including the market excess returns. We then standardised the yield curve factors with a mean of zero and a standard deviation of one, which simplifies parameter interpretation.

In addition, we divide our dataset into two segments, before and after the 1st of January 1985. In Chapter 2, we highlighted several studies which documented that the predictive power of the slope had diminished since the mid-1980s. Existing research, paired with the empirical evidence from Figure 3.5.9, which highlights a transition from a long-term upward interest rate cycle to a downward cycle in the mid-1980s, motivated us to divide the sample at the 1985 juncture.

Lastly, we calculate the cumulative 12-month and 6-month log excess returns as the sum of market excess log returns.

3.4 Data Summary and Descriptive Statistics

The response variable of interest is the cumulative, overlapping 12-month and 6-month logarithmic excess return for the entire sample period, pre-1985 period, and post-1985 period. The construction of covariates, the level, slope, and curvature, is described in section 3.2. The following three figures highlight the descriptive statistics of the total logarithmic excess returns conditioned on different quartiles and the full distribution for all three yield curve factors.

Level									
12-months	***Full Sample***			***Pre-1985***			***Post-1985***		
	mean	**st.dev**	**skewness**	**mean**	**st.dev**	**skewness**	**mean**	**st.dev**	**skewness**
Full Distribution	5,54 %	16,55 %	-0,85	2,17 %	16,61 %	-0,58	7,35 %	16,37 %	-1,02
1st Quartile	11,93 %	14,80 %	-0,95	6,96 %	9,34 %	-0,54	13,55 %	14,15 %	-0,30
2nd & 3rd Quartile	2,09 %	16,43 %	-1,01	-1,21 %	17,93 %	-0,64	3,66 %	16,94 %	-1,30
4th Quartile	6,03 %	16,49 %	-0,58	4,15 %	17,98 %	0,09	8,54 %	15,19 %	-0,80

Table 3.4.1*: 12-month cumulative log excess return statistics. The figure reports the mean, the standard deviation, and the standardised third moment (skewness) for the full sample, the pre-1985 sample, and the post-1985 sample, conditional on different quartiles of the level factor, as well as for the full distribution.*

Level									
6-months	***Full Sample***			***Pre-1985***			***Post-1985***		
	mean	**st.dev**	**skewness**	**mean**	**st.dev**	**skewness**	**mean**	**st.dev**	**skewness**
Full Distribution	2,68 %	11,75 %	-0,79	0,93 %	12,08 %	-0,23	3,63 %	11,48 %	-1,15
1st Quartile	5,97 %	10,76 %	-0,90	2,38 %	9,10 %	-0,95	7,71 %	9,61 %	-0,31
2nd & 3rd Quartile	1,40 %	10,82 %	-1,13	-0,32 %	13,21 %	-0,24	1,17 %	11,93 %	-1,48
4th Quartile	1,96 %	13,69 %	-0,38	1,96 %	12,07 %	0,42	4,47 %	10,99 %	-0,77

Table 3.4.2*: 6-month cumulative log excess return statistics. The figure reports the mean, the standard deviation, and the standardised third moment (skewness) for the full sample, the pre-1985 sample, and the post-1985 sample, conditional on different quartiles of the level factor, as well as for the full distribution.*

Tables 3.4.1 and 3.4.2 show that the level of the yield curve impacts the excess returns (reported as mean), with lower average yields (1st quartile) resulting in higher excess returns compared to higher average yields (4th quartile). The inverse relationship holds across all sub-samples and for both return horizons.

The volatility, measured as the standard deviation of cumulative excess returns, decreases when average yields are low compared to a high average yield. This is particularly evident in the pre-1985 sample, where volatility decreases from 17.98% to 9.34% for the 12-month holding period.

Skewness, measured as the third standardised moment, reduces for the full sample, and becomes positive in the pre-1985 sample as the average interest rate level becomes high. Remarkably, this trend shifts in the post-1985 sample, where the returns become more negatively skewed as interest rates increase. Nonetheless, this measure of skewness is susceptible to large sampling errors, particularly in distributions with heavy tails, and it is vulnerable to outliers (Giordani & Halling, 2019).

Slope									
12-months	***Full Sample***			***Pre-1985***			***Post-1985***		
	mean	**st.dev**	**skewness**	**mean**	**st.dev**	**skewness**	**mean**	**st.dev**	**skewness**
Full Distribution	5,54 %	16,55 %	-0,85	2,17 %	16,61 %	-0,58	7,35 %	16,37 %	-1,02
1st Quartile	-2,19 %	18,48 %	-0,42	-6,60 %	20,58 %	-0,22	4,37 %	18,55 %	-0,60
2nd & 3rd Quartile	7,79 %	15,89 %	-1,05	5,77 %	14,41 %	-0,64	7,70 %	16,57 %	-1,21
4th Quartile	8,77 %	12,99 %	-0,48	3,79 %	12,59 %	0,87	9,65 %	12,79 %	-0,82

Table 3.4.3: *12-month cumulative log excess return statistics. The figure reports the mean, the standard deviation, and the standardised third moment (skewness) for the full sample, the pre-1985 sample, and the post-1985 sample, conditional on different quartiles of the slope factor, as well as for the full distribution.*

Slope									
6-months	***Full Sample***			***Pre-1985***			***Post-1985***		
	mean	**st.dev**	**skewness**	**mean**	**st.dev**	**skewness**	**mean**	**st.dev**	**skewness**
Full Distribution	2,68 %	11,75 %	-0,79	0,93 %	12,08 %	-0,23	3,63 %	11,48 %	-1,15
1st Quartile	-1,49 %	12,28 %	-0,25	-5,14 %	14,31 %	-0,02	2,22 %	11,15 %	-0,12
2nd & 3rd Quartile	4,41 %	11,15 %	-1,22	3,33 %	10,72 %	-0,14	4,55 %	11,57 %	-1,85
4th Quartile	3,40 %	11,35 %	-0,60	2,22 %	9,94 %	0,72	3,22 %	11,45 %	-0,73

Table 3.4.4: *6-month cumulative log excess return statistics. The figure reports the mean, the standard deviation, and the standardised third moment (skewness) for the full sample, the pre-1985 sample, and the post-1985 sample, conditional on different quartiles of the slope factor, as well as for the full distribution.*

Summary statistics in Table 3.4.3 and 3.4.4 display a consistent increase in excess returns (mean) across all samples and horizons during periods of an upward-sloping yield curve (4th quartile), contrasted with periods when the yield curve is either narrow or inverted (1st quartile). This trend is particularly pronounced in the full sample and pre-1985 sample.

Conversely, an inverse relationship is observed regarding volatility. The volatility across all samples and horizons decreases during periods of upward-sloping term structure compared to periods when the yield curve is narrow or inverted.

Skewness becomes increasingly negative for the full sample as term structure steepens but shows contrasting behaviour between the two sub-samples. For the pre-1985 sample, skewness turns negative to positive as the yield curve steepens, while for the post-1985 sample, the skewness becomes increasingly negative as the term structure steepens.

12-months	Curvature								
	Full Sample			*Pre-1985*			*Post-1985*		
	mean	st.dev	skewness	mean	st.dev	skewness	mean	st.dev	skewness
Full Distribution	5,54 %	16,55 %	-0,85	2,17 %	16,61 %	-0,58	7,35 %	16,37 %	-1,02
1st Quartile	3,25 %	21,05 %	-0,81	-4,16 %	19,84 %	-0,49	6,78 %	20,63 %	-1,05
2nd & 3rd Quartile	5,82 %	15,11 %	-0,78	3,15 %	14,79 %	-0,67	7,55 %	15,14 %	-0,88
4th Quartile	7,27 %	13,59 %	0,27	6,79 %	14,22 %	0,51	7,53 %	13,68 %	-0,88

Table 3.4.5*: 12-month cumulative log excess return statistics. The figure reports the mean, the standard deviation, and the standardised third moment (skewness) for the full sample, the pre-1985 sample, and the post-1985 sample, conditional on different quartiles of the curvature factor, as well as for the full distribution.*

6-months	Curvature								
	Full Sample			*Pre-1985*			*Post-1985*		
	mean	st.dev	skewness	mean	st.dev	skewness	mean	st.dev	skewness
Full Distribution	2,68 %	11,75 %	-0,79	0,93 %	12,08 %	-0,23	3,63 %	11,48 %	-1,15
1st Quartile	3,17 %	11,97 %	-0,53	-1,01 %	13,12 %	-0,34	5,32 %	10,74 %	-0,48
2nd & 3rd Quartile	2,59 %	11,81 %	-1,06	0,80 %	11,78 %	-0,34	3,80 %	11,50 %	-1,47
4th Quartile	2,37 %	11,39 %	-0,48	3,13 %	11,16 %	0,39	1,62 %	11,84 %	-1,05

Table 3.4.6*: 6-month cumulative log excess return statistics. The figure reports the mean, the standard deviation, and the standardised third moment (skewness) for the full sample, the pre-1985 sample, and the post-1985 sample, conditional on different quartiles of the curvature factor, as well as for the full distribution.*

Table 3.4.5 shows 12-month cumulative excess returns conditioned on curvature. We note an increase in the mean for all three samples for high curvature readings, representing a concave term structure (4th quartile) compared to a convex term structure (1st quartile). Table 3.4.6 show that the relationship is less concise for the 6-month holding period. We observe an increase in the mean for the full sample and pre-1985 sample, while it decreases with curvature in the post-1985 period.

The volatility decreases across all sub-samples for the 12-month holding period when curvature is high. The mean and volatility demonstrate properties similar to the slope. Again, the results for the 6-month holding period are less concise, displaying marginal changes in volatility in both directions.

For the full and pre-1985 samples, the return distribution displays negative skewness under conditions of low curvature, contrasting with positive skewness observed in periods of high curvature. Interestingly, while the post-1985 sample follows a similar pattern of diminishing skewness during periods of high curvature, it does not transition into positive territory. For the 6-month skewness, the difference between the pre-and post-1985 samples are more pronounced, with skewness decreasing in curvature in the pre-1985 sample and increasing in the post-1985 period.

3.5 Preliminary Analysis

This section provides further preliminary analysis beyond the descriptive statistics presented in section 3.4. First, we visualise and describe the characteristics of the return series, which serve as the response variable in the analysis. We then highlight why an Ordinary Least Squares (OLS) model is insufficient to capture the dynamics for our data, and lastly, we describe the yield curve factors.

3.5.1: Cumulative and Conditional Returns

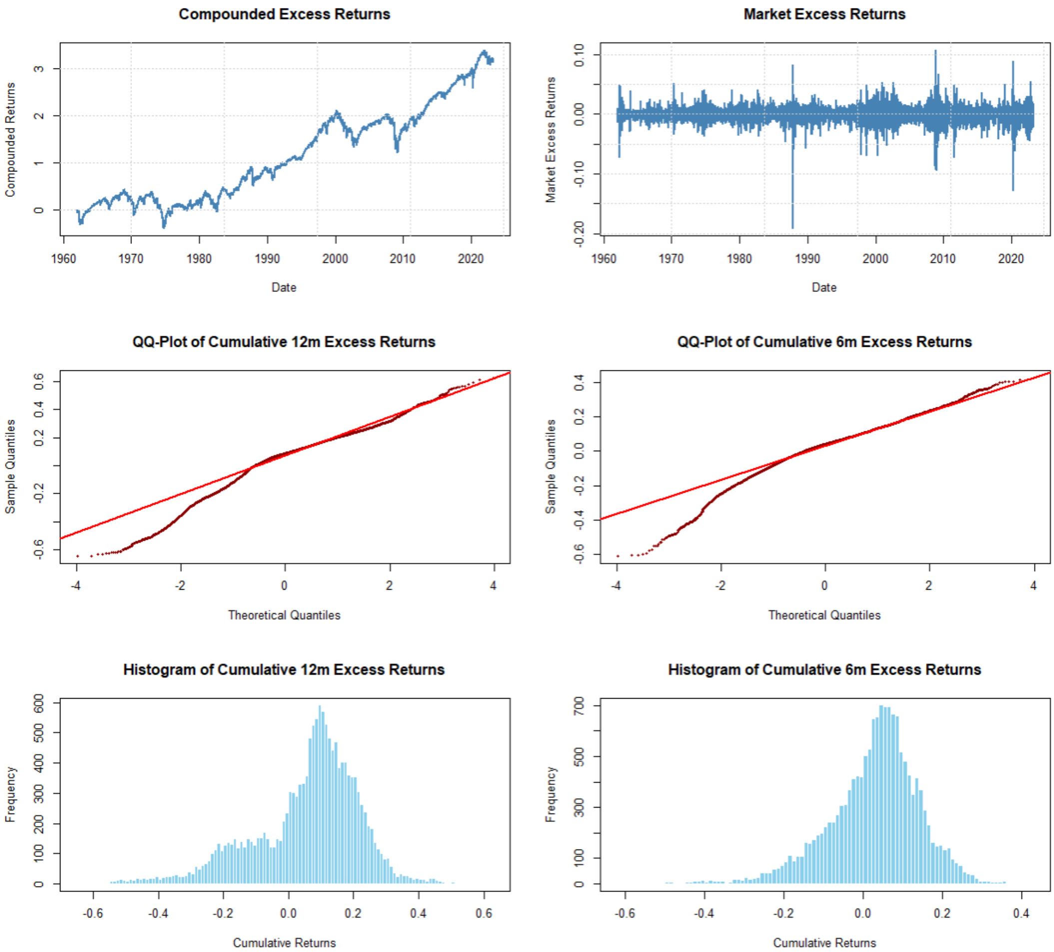

Figure 3.5.1: *Excess Return Characteristics. The top left picture shows the compounded daily log returns for the entire sample period, starting in 1962. The top right picture highlights daily market excess log returns. The middle-left (right) picture shows the QQ-plot of sample quantiles vs theoretical quantiles for cumulative 12-month (6-month) excess returns. The bottom left (right) picture shows the histogram of cumulative 12-month (6-month) excess log returns.*

The top left image in Figure 3.5.1 displays the compounded market excess returns from the inception of our sample in 1962. A quick look reveals periods of drawdowns in the compounded returns. The image on the top right focuses on daily excess returns plotted over time. This visualises the evidence of volatility clustering and draws attention to extreme return values, particularly the downside.

The two middle images present QQ-plots comparing theoretical and sample quantiles of 12- and 6-month cumulative excess returns. If our return distribution perfectly followed a normal distribution, the data points in these plots would align precisely with the bright red straight line. However, the noticeable deviations from this line suggest that our cumulative return series do not follow the normal distribution, specifically in the left tail region.

Lastly, the two bottom plots in Figure 3.5.1 provide further evidence of non-normality. Here, we can see signs of skewness and heavier tails in the return distribution. These findings highlight one of the stylised properties of univariate financial time series, namely that tails are much thicker than the Gaussian.

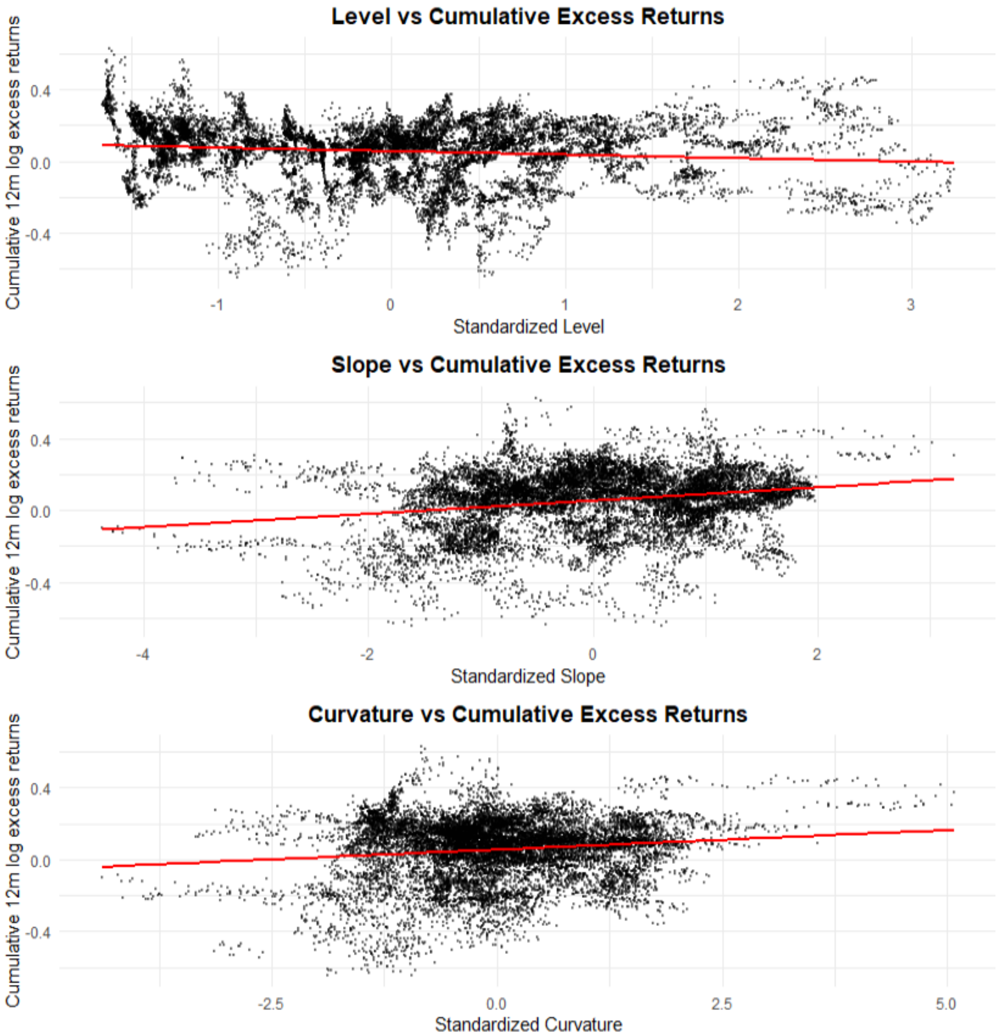

Figure 3.5.2*: Full samples; scatter plots of standardised yield curve factors (x-axis) vs cumulative 12-month excess log returns (y-axis) with a fitted simple linear regression line. The top plot shows the standardised level vs returns, the middle plot shows the standardised slope vs returns, and the bottom plot shows the standardised curvature vs returns. Scatter plots for the cumulative 6-month excess log returns are provided in Figure A.3.1 in the appendix.*

Figure 3.5.2 displays scatter plots of the standardised yield curve factors on the x-axis against the cumulative 12-month excess log returns on the y-axis. Figure A.3.1 in the appendix displays the same scatter plots for the 6-month horizon. A simple regression line is fitted to each scatter plot. These plots highlight subtle relationships between the response variable and our covariates.

The regression line displays a negative relationship for the level factor while exhibiting a positive relationship for the slope and curvature factors. Additionally, we observe distinct clusters of data points for the slope and curvature factors and several outliers for each of the three factors. Similar observations can be seen for cumulative 6-month excess log returns, as shown in Figure A.3.1, except for the curvature factor, where the regression line appears to flatten out.

3.5.2: Ordinary Least Squares Analysis

We extend our preliminary analysis by estimating an OLS regression for the cumulative 12- and 6-month log excess returns on the yield curve factors. Although the OLS analysis is not the primary objective of this thesis, we report the full sample results for context. It is worth noting that the limitations of the OLS model, as we will demonstrate, apply to both the full sample and the sub-samples.

	Level			***Slope***			***Curvature***		
	Estimate	**Std. Error**	**T-Value**	**Estimate**	**Std. Error**	**T-Value**	**Estimate**	**Std. Error**	**T-Value**
Intercept	0,0554	0,0013	41,49	0,0554	0,0013	42,29	0,0554	0,0013	41,57
Beta	-0,0195	0,0013	-14,60	0,0372	0,0013	28,42	0,0218	0,0013	16,33

Table 3.5.3*: OLS regression results, cumulative 12-month excess returns on yield curve factors. We report the estimated value, the standard error, and the t-statistic for both the intercept and the beta coefficient for each yield curve factor. All variables are in logs.*

	Level			***Slope***			***Curvature***		
	Estimate	**Std. Error**	**T-Value**	**Estimate**	**Std. Error**	**T-Value**	**Estimate**	**Std. Error**	**T-Value**
Intercept	0,0268	0,0009	28,37	0,0268	0,0009	28,46	0,0268	0,0010	28,22
Beta	-0,0122	0,0009	-12,96	0,0153	0,0009	16,29	0,0007	0,0010	0,76

Table 3.5.4*: OLS regression results, cumulative 6-month excess returns on yield curve factors. We report the estimated value, the standard error, and the t-statistic for both the intercept and the beta coefficient for each yield curve factor. All variables are in logs.*

The results in Tables 3.5.3 and 3.5.4 illustrate how cumulative 12- and 6-month excess returns react to one-standard deviation shifts in the yield curve factors. Specifically, when the level increases by one standard deviation from the mean, the cumulative 12-month excess returns decrease by -1.95%, while the cumulative 6-month excess returns decrease by -1.22%. Contrarily, as implied by the OLS model, if the slope or curvature factor rises by one standard deviation from the mean, the cumulative 12-month excess returns increase by 3.72% and 2.18%, respectively (1.53% and 0.7% for the 6-month horizon, but significance drastically falls for the curvature).

However, the OLS regression model makes several key assumptions about the errors which underpin the properties and reliability of the model estimates which we summarise in Table 3.5.5 below.

(1) $E(u_t) = 0$	Average value of errors is zero
(2) $var(u_t) = \sigma^2 < \infty$	Variance of errors is constant
(3) $cov(u_i, u_j) = 0$	Errors are linearly independent
(4) $cov(u_t, x_t) = 0$	Covariates are non-stochastic
(5) $u_t \sim N(0, \sigma^2)$	Errors are normally distributed

Table 3.5.5: *The table above highlights the mathematical notation and interpretation of OLS assumptions. (Brooks, 2019)*

Various tests can be performed to handle the violation of the assumptions outlined in the table above. Some examples are White's test for heteroskedasticity, The Durbin-Watson test for autocorrelation, and the Bera-Jarque test for normality. While a deep dive into all assumptions is not the focus of this thesis, we will tackle the most critical assumption violation, especially given the nature of our data set.

Our response variable, cumulative excess log returns, is constructed using annual and semi-annual overlapping data with daily returns. This results in massive autocorrelation, making any test for this somewhat redundant. If we ignore autocorrelation, the coefficient estimates produced by OLS will still be unbiased, but inefficient, and standard errors will be incorrect (Brooks, 2019). We present the OLS results with heteroskedasticity- and autocorrelation-consistent (HAC) standard errors in Tables 3.5.6 and 3.5.7 below.

	Level			*Slope*			*Curvature*		
	Estimate	**Std. Error**	**T-Value**	**Estimate**	**Std. Error**	**T-Value**	**Estimate**	**Std. Error**	**T-Value**
Intercept	0,0554	0,0165	3,35	0,0554	0,0170	3,25	0,0554	0,0191	2,90
Beta	-0,0195	0,0130	-1,50	0,0372	0,0132	2,81	0,0218	0,0173	1,26

Table 3.5.6*: OLS regression results, cumulative 12-month excess returns on yield curve factors. We report the estimated value, the HAC standard error, and the t-statistic for the intercept and beta coefficients for each yield curve factor.*

	Level			Slope			Curvature		
	Estimate	Std. Error	T-Value	Estimate	Std. Error	T-Value	Estimate	Std. Error	T-Value
Intercept	0,0268	0,0090	2,99	0,0268	0,0096	2,79	0,0268	0,0103	2,60
Beta	-0,0122	0,0066	-1,87	0,0153	0,0093	1,66	0,0007	0,0094	0,08

Table 3.5.7*: OLS regression results, cumulative 6-month excess returns on yield curve factors. We report the estimated value, the HAC standard error, and the t-statistic for the intercept and beta coefficients for each yield curve factor.*

From the above results, we see a substantial decrease in significance compared to the findings in Tables 3.5.3 and 3.5.4 once we account for heteroskedasticity and autocorrelation. Notably, the slope factor emerges as the most significant covariate in predicting excess returns, as indicated by the largest t-statistic.

Figure 3.5.1 showed that our dataset contained outliers, which is common in financial return data. Because OLS aims to minimise the sum of the squared residuals, an outlier can disproportionately influence the line of best fit. This is because squaring the residuals amplifies the distances of points far from the line, so the line will be more 'pulled' towards an outlier.

We utilise a simple boxplot to highlight outliers in our cumulative excess returns. This is a visual representation of the five-number summary of a dataset: the minimum, the first quartile, the median, the third quartile, and the maximum.

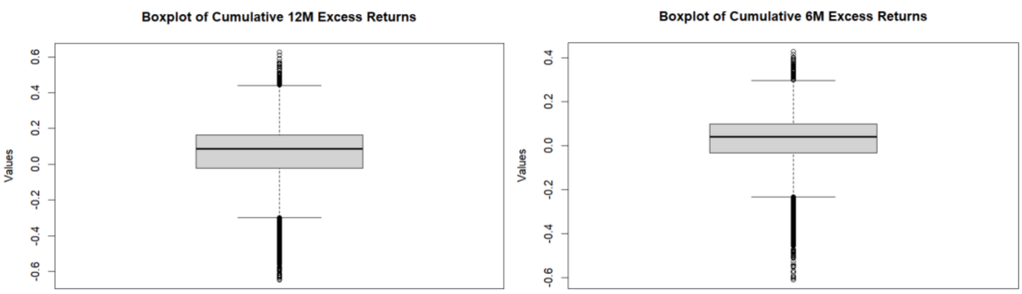

Figure 3.5.8: *Boxplot highlighting outliers of cumulative 12- and 6-month excess log returns. The box indicates the middle 50% of values, the line inside represents the median, and points outside the whiskers are potential outliers. It provides a quick way to spot outliers and understand the data's spread.*

Figure 3.5.8 shows that our response variable has several outliers, highlighted as the data points that fall outside the "whiskers". The "whiskers" represent the data outside the middle 50%.

In summary, our OLS analysis displays some evidence of a relationship between the yield curve factors, particularly the slope, and the mean of cumulative excess returns. However, our data violates some of the assumptions underlying the standard linear regression and contains several outliers, which leads to unreliable model estimates. Given these observed complications, it becomes apparent that a more complex and flexible modelling approach, such as GAMLSS, is warranted. Moreover, our analytical interests extend beyond merely exploring the mean behaviour. We aim to delve into the shape of the distribution, thereby necessitating the adoption of a modelling technique like GAMLSS capable of modelling multiple parameters of the distribution of the response variable.

3.5.3: Yield Curve Factors

This section examines our covariates, level, slope, and curvature, using time-series plots and histograms to gain insights into their dynamics over time.

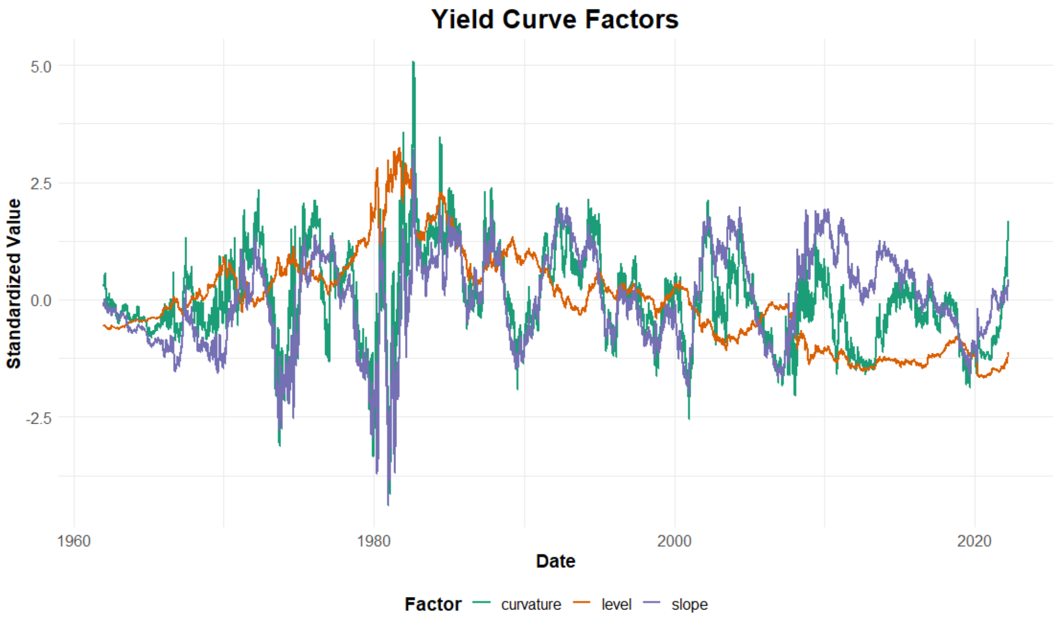

Figure 3.5.9*: Time series of standardised values of yield curve factors for the full sample. The orange line depicts the standardised level factor, the purple line shows the standardised slope factor, and the green line highlights the standardised curvature factor. Time series for subsamples are provided in Figures A.3.2 and A.3.3 in the appendix.*

From Figure 3.5.9, three significant observations warrant attention. Firstly, examining the level factor, a clear shift in interest rate environments is observed. Prior to 1985, the trend indicates a period of steadily rising long-term interest rates. However, post-1985 marks a pivot in this trend, indicating a shift towards a prolonged environment of decreasing interest rates.

Secondly, a correlation can be observed between the curvature and slope factors. For the full sample, the correlation stands at 72% between these two factors. However, a more detailed inspection reveals that this relationship varies over time. Specifically, the correlation is 88% in the pre-1985 sample, while it moderates to 70% in the post-1985 sample.

Lastly, caution should be exercised using the level factor for prediction. Since our model is built to predict cumulative excess returns conditional on the level factor's deviation from the mean, the prediction relies on a limited number of observations. This scarcity of data points could compromise the reliability of the results.

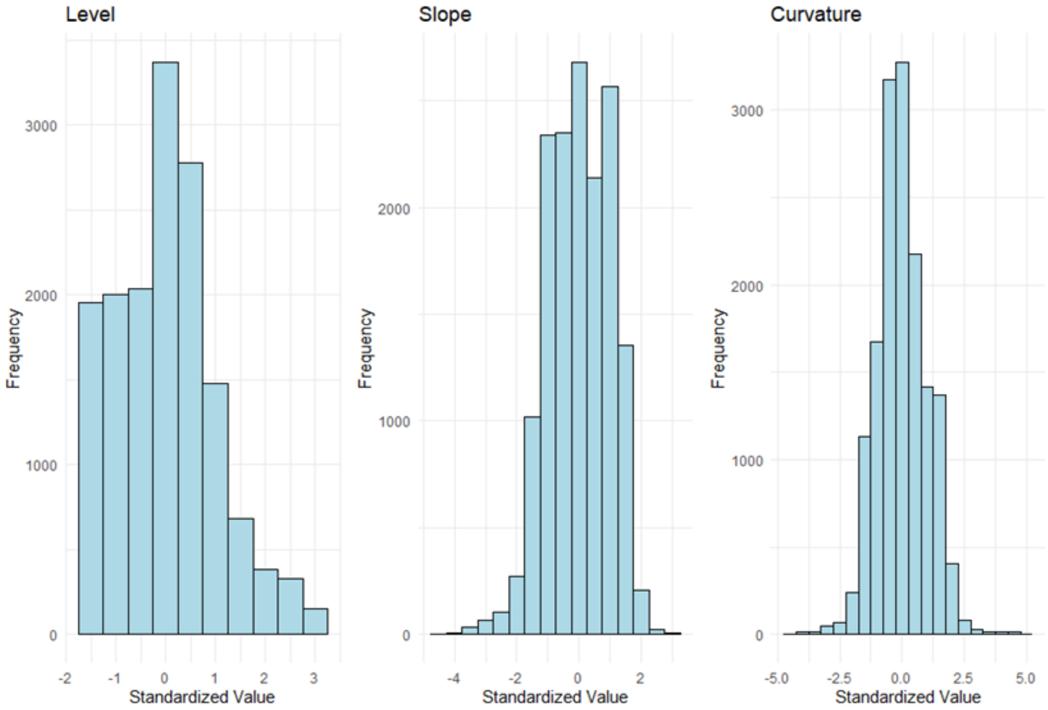

Figure 3.5.10*: This figure presents histograms of the standardised level, slope, and curvature factors for the full sample. Each histogram visually illustrates the distribution of its respective yield curve factor. Figures A.3.4 and A.3.5 in the appendix provide histograms of yield curve factors for the subsamples.*

Figure 3.5.10 presents histograms of standardised yield curve factors. Notably, the level factor values always stay above two negative standard deviations. This is an important implication when modelling conditional return distributions based on deviations from the mean. Both the slope and curvature histograms exhibit fatter tails, with the slope also displaying pronounced negative skewness. The evidence gathered from these histograms is important in understanding the underlying data structure and ensuring the robustness of our model.

4. Methodology

In this chapter, we lay the groundwork for our analysis presented in Chapter 5 by providing a thorough overview of the fundamentals and assumptions that underpin our modelling approach. We start by introducing the GAMLSS framework and its advantages. We then propose the preferred distributions and describe the models used in the primary analysis. In the final section, we provide an overview of model estimation techniques, model selection criteria and the computational approach for risk metrics.

4.1: Overview of GAMLSS model

The main analysis in this thesis is based on the Generalized Additive Model for Location Scale and Shape (GAMLSS). Introduced by Rigby and Stasinopoulos (2005), GAMLSS builds upon two earlier models. The first is the Generalized Linear Models (GLMs), proposed by Nelder and Wedderburn (1972), and the second is the Generalized Additive Models (GAMs), developed by Hastie and Tibshirani (1990). The GAMLSS method extends their ideas, providing a more flexible and comprehensive tool for our analysis.

Model	Characteristics	Function
OLS	Linear function for the mean. Higher moments are constant.	$y \sim N(\mu, \sigma^2)$ 4.1 $\boldsymbol{\mu} = \boldsymbol{X\beta}$
GLM	Link function allowing a linear function to have a non-linear relationship with y.	$y \sim ExpFamily(\mu, \phi)$ 4.2 $\boldsymbol{g}(\boldsymbol{\mu}) = \boldsymbol{X\beta}$
GAM	Non-parametric function for the response variable.	$y \sim ExpFamily(\mu, \phi)$ 4.3 $\boldsymbol{g}(\boldsymbol{\mu}) = \boldsymbol{X\beta} + s(x_1) + \cdots + s_j(x_j)$

GAMLSS has two main advantages for modelling the impact of covariates on distribution parameters. Firstly, it provides flexibility in choosing distributional assumptions. Unlike GLM or GAM, which are limited to normal or exponential family distributions, GAMLSS allows us to select from a wide range of distributions that better represent our underlying data-generating process. A GAMLSS model is expressed as $\boldsymbol{y} \sim \mathcal{D}(\boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{\nu}, \boldsymbol{\tau})$, where $\mathcal{D}$ represents the distribution of the response variable y. This general distribution can take any form with up to four parameters.

Secondly, GAMLSS enables us to model any distribution parameter as a function of the predictive variables. This means we can examine the impact of the yield curve level factors, the predictive variables, on the entire shape of the return distribution, not just the expected value. By considering the relationship between the covariates and the distribution parameters, we gain insights into how different factors influence the variance, skewness, and tails of the response variable's distribution. The parameters of the distribution can be modelled as a function of the explanatory variables as follows:

$$\boldsymbol{y} \sim \mathcal{D}(\boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{\nu}, \boldsymbol{\tau}) \qquad 4.4$$
$$g_1(\boldsymbol{\mu}) = \boldsymbol{X_1\beta_1} + s_{11}(x_{11}) + \cdots + s_{1J_1}(x_{1J_1})$$
$$g_2(\boldsymbol{\sigma}) = \boldsymbol{X_2\beta_2} + s_{21}(x_{21}) + \cdots + s_{2J_2}(x_{2J_2})$$
$$g_3(\boldsymbol{\nu}) = \boldsymbol{X_3\beta_3} + s_{31}(x_{31}) + \cdots + s_{3J_3}(x_{3J_3})$$
$$g_4(\boldsymbol{\tau}) = \boldsymbol{X_4\beta_4} + s_{41}(x_{41}) + \cdots + s_{4J_4}(x_{4J_4})$$

In the model, β_j represents the linear *parametric* estimator of $g(.)$, while $s(.)$ denotes a smoothing *non-parametric* function applied to some continuous explanatory variables. The **gamlss** package offers various non-parametric functions that can be used to improve the model's fit by applying them to the explanatory variables. However, in this thesis, our primary objective is interpretation rather than maximising the model's predictive power. Therefore, we will only utilise the parametric estimators β_j, as the non-parametric estimators are not easily interpretable, and we cannot perform significance tests in the same manner as with linear models.

4.2: Modelling approach

Our goal is to test the hypothesis of whether yield curve factors influence the shape of the conditional distribution of equity returns. Specifically, we want to check if the yield curve factors influence the asymmetry of the return distribution and its risk-return characteristics. To do that in the most parsimonious and easily interpretable way, we utilise a linear model with different distributional assumptions of the error terms. Gradually expanding the model flexibility allows us to isolate the effects of our predictors on the different distribution parameters. More specifically, we use the "Conditional-Skew-T" model to predict the time-varying asymmetry of the distribution and compare the results to three other benchmark models, namely the “Normal model”, “Symmetric-T model”, and “Constant-Skew-T model”. As a formal model comparison, we use the GAIC model selection criteria. This approach follows the logic in Fernandez and Steel (1998) that introduces skewness to the symmetric t-distribution. However, the more specific approach to testing the distribution parameters on equity returns follows the methodology outlined in Giordani and Halling (2019). The latter is more convenient for our purpose, as the specific models proposed are directly applicable to analyse the shape of the return distribution.

Equity returns often exhibit several key characteristics that require more flexible distribution families to accurately represent the underlying data-generating process. As noted in Chapter 3, our dataset displays asymmetry and thicker tails than a normal distribution. If the error terms are not normally distributed but, for example, have a heavy-tailed distribution, OLS estimates of regression coefficients can be disproportionately affected by outliers or extreme values. This makes OLS model estimates less efficient, increasing the model variance.

To tackle the problem of asymmetry and heavy tails in our data, we employ a t-distribution, which allows for heavier tails than a Gaussian, and a skew t-distribution, which introduces asymmetry. A more flexible distribution for error terms can make estimates more robust against non-normality since more flexible distribution parameters can now capture outliers. Another method to deal with outliers is to move away from a linear estimator and use splines or other non-linear methods. The problem with non-linearity is complex interpretation, and in **gamlss**, the standard errors must be treated with caution (Rigby et al., 2017). Therefore, we opt for the linear model, but by changing the distributional assumption, we mitigate the effect of outliers and leverage points.

Another important issue in our dataset is the presence of overlapping observations. Since we are forecasting the cumulative excess returns a year ahead, $rx = (p_{t+252} - p_t)$, we have 251 overlapping observations, and the same logic applies for the semi-annual sample. The autocorrelation introduced by this large number of observations will significantly overestimate the t-statistics. In an OLS model, one would employ autocorrelation consistent standard errors. However, the **gamlss** package lacks any means of adjusting for overlapping observations. Therefore, we employ a solution proposed by Giordani and Halling (2019) to adjust the t-statistics and log-likelihood, described in detail in Sections 4.4 and 4.5.

4.3: Distribution Family and Model Specifications

This chapter presents the models used in the analysis, including distributional assumptions and parameter functions. To methodically document the marginal improvements of gradually increasing model flexibility, we start with the Normal distribution assumption and move gradually to distributions with more parameters.

Since we are using a simple linear model and only change the distributional assumptions of the error terms, we can interpret the $\beta_{i,\phi}$ coefficients similarly to OLS coefficients. After adjusting for autocorrelation, we can also perform tests of statistical significance for the linear coefficients.

In all models, we are using standardised log values for the covariates. This implies that the $\beta_{i,\phi}$ coefficients represent the percentage change in the response variable, given one standard deviation change in the yield curve factors.

Normal model

Our Normal model is equivalent to a linear model estimated by OLS since we use a normal distribution and only model the mean parameter as a function of our covariates. The Normal model has the following functional form:

$$\begin{aligned} &\boldsymbol{y} \sim N(\boldsymbol{\mu}, \boldsymbol{\sigma}) \\ &g_1(\boldsymbol{\mu}) = \beta_{0,\mu} + \beta_{1,\mu} x_t \\ &g_2(\boldsymbol{\sigma}) = \beta_{0,\sigma} \end{aligned} \tag{4.5}$$

If the data represents thicker-than-normal tails, a Normal model with normal tail behaviour needs to increase the scale (σ) to capture observations in the tails (Fernandez & Steel, 1998). This characteristic will misrepresent the relevant risk metrics if the distribution is misspecified.

Symmetric-T model

The Symmetric-T model leverages the Student's t-distribution, allowing for longer-than-normal tails. This allows for more robust statistical modelling of datasets with outliers and fat-tail errors. The t-distribution, defined by three parameters $y_{t,t+h} \sim t(\mu, \sigma, \nu)$, extends the typical normal distribution used in linear models. Its degree of freedom parameter ν aids in reliable mean estimation (Lange et al., 2019). Unlike models with normal error assumptions, which are vulnerable to outliers, the heavy tails of the t-distribution result in less outlier sensitivity and more robust mean estimates. In a linear regression context, using a t-distribution can result in estimates that are less sensitive to changes in a small number of data points.

We construct the Symmetric-T model with a time-dependent mean, while keeping the dispersion and degrees of freedom constant. This way, we can estimate the significance of the predictive variables on the mean, while controlling for outliers in the sample. The Symmetric-T model in functional form is written as follows:

$$\begin{aligned} &\boldsymbol{y} \sim t(\boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{\nu}) \qquad (4.6) \\ &g_1(\boldsymbol{\mu}) = \beta_{0,\mu} + \beta_{1,\mu} x_t \\ &g_2(\boldsymbol{\sigma}) = \beta_{0,\sigma} \\ &g_3(\boldsymbol{\nu}) = \beta_{0,\nu} \end{aligned}$$

Interpretation of the model is still straightforward. The $\boldsymbol{\mu}$ is the mode, but since the distribution is symmetric, the mode equals the mean. $\boldsymbol{\sigma}$ is the standard deviation, while $\boldsymbol{\nu}$ is the degrees of freedom parameter that controls the length of the tails. When the degrees of freedom are high, the t-distribution is similar to a normal distribution. However, when the degrees of freedom are low, the t-distribution has fatter tails. Note that the degrees of freedom parameter, ν, in the Symmetric-T model is equivalent to τ in the following two models, while ν in the following two models represents the asymmetry parameter.

Constant-Skew-T model

In our Constant-Skew-T model, we employ the Skewed Student-t Type 3 (ST3) distribution created by Fernandez and Steel (1998). The ST3 distribution adds another layer to the t-distribution by introducing a shape parameter ν that controls the asymmetry of the distribution. This method allows linear models to account for skewed error distributions with fat tails. By including a separate shape parameter ν, in addition to the tail parameter τ, the Constant-Skew-T mitigates the risk of the model treating asymmetry as excess kurtosis (Giordani & Halling, 2019).

The shape parameter is the main output in the analysis and acts as a proxy for the skewness of the distribution. As previously mentioned, the skewness defined as a third moment is a volatile metric highly affected by outliers. An ST3 distribution is designed so that the shape parameter, ν, controls the allocation of probability mass to each side of the mode as:

$$\frac{P(y_t \geq m|\nu_t)}{P(y_t \leq m|\nu_t)} = \nu_t^2 \qquad 4.7$$

The asymmetry parameter is defined $0 < \boldsymbol{\nu} < \infty$, where $\boldsymbol{\nu} = 1$ is a symmetric distribution. A major advantage of using the ST3 distribution is that the approach completely separates the effect of the skewness parameter, ν, and the tail parameter, τ, facilitating their interpretation and making prior independence between the two a plausible assumption (Fernandez & Steel, 1998). In a functional form, our Constant-Skew-T model is presented in the following way:

$$\begin{aligned} &\boldsymbol{y} \sim skewt(\boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{\nu}, \boldsymbol{\tau}) \qquad 4.8 \\ &g_1(\boldsymbol{\mu}) = \beta_{0,\mu} + \beta_{1,\mu} x_t \\ &g_2(\boldsymbol{\sigma}) = \beta_{0,\sigma} \\ &g_3(\boldsymbol{\nu}) = \beta_{0,\nu} \\ &g_4(\boldsymbol{\tau}) = \beta_{0,\tau} \end{aligned}$$

The interpretation of the ST3 is more complicated than for the regular Student's t-distribution since the statistical moments are generally a combination of all four distribution parameters. However, since the skewness and tail parameters are separated, we can infer the effect of the covariates on the shape of returns.

Parameter $\boldsymbol{\tau}$ now controls the degrees of freedom and has the same interpretation as in the Symmetric-T model (stated as ν in the Symmetric-T). $\boldsymbol{\mu}$ is, in general, the mode for all values when $\boldsymbol{\nu} \neq \mathbf{1}$, with the special case of symmetry, then the mode becomes the mean. $\boldsymbol{\sigma}$ now becomes the dispersion parameter.

Conditional-Skew-T

We introduce the Conditional-Skew-T model to effectively isolate and model the time-varying effect of the shape parameter ν. This model differs from the Constant-Skew-T by modelling the shape parameter as a function of our covariates. While the GAMLSS framework theoretically enables us to also model the kurtosis or degrees of freedom as a function of covariates, the limited number of extreme observations presents a challenge. Therefore, attempting to model the thickness of the tails as a function of covariates could lead to high model variance and unstable parameters.

As a compromise, we adopt a distribution family that includes a tail parameter but treats it as a constant in our modelling approach. This decision allows us to ensure model stability while separating the effect of tail behaviour and asymmetry on the shape of the distribution. The Conditional-Skew-T model has following functional form:

$$\begin{aligned} &\boldsymbol{y} \sim skewt(\boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{\nu}, \boldsymbol{\tau}) \\ &g_1(\boldsymbol{\mu}) = \beta_{0,\mu} \\ &g_2(\boldsymbol{\sigma}) = \beta_{0,\sigma} \\ &g_3(\boldsymbol{\nu}) = \beta_{0,\nu} + \beta_{1,\nu} \\ &g_4(\boldsymbol{\tau}) = \beta_{0,\tau} \end{aligned} \qquad 4.9$$

The interpretation is similar to the Constant-Skew-T model, but now we also estimate the linear $\beta_{1,\nu}$ coefficient. The coefficient measures the time-varying effect of the yield curve factors on the shape of cumulative equity returns and is crucial for our analysis.

4.4: Model Estimation Techniques

Estimation in a linear parametric GAMLSS model is done by maximising the log-likelihood ℓ given by:

$$\ell = \sum_{i=1}^{n} \log f(y_i|\mu_i, \sigma_i, \nu_i, \tau_i) \qquad 4.10$$

Where $f(.)$ represents the Probability Density Function (PDF) of the response variable. While a more general, non-parametric model is estimated using maximum penalised log-likelihood estimation, we focus only on parametric models; thus, the simple maximum log-likelihood algorithm holds. It is also important to note that the log-likelihood function assumes that the observations in the response variables are independent. Clearly, the independence assumption is violated due to volatility clustering and autocorrelation present in our data. Violating the independence assumption leads to incorrect estimates of the standard errors, which makes the t-statistics incorrect. Therefore, a correction for autocorrelation should be made to compute standard errors and t-statistics. In our particular case of 251 overlapping observations for the 12-month holding period, the t-statistics are significantly overestimated. The same reasoning applies to other holding periods. Because, to the best of our knowledge, the **gamlss** package lacks the means of correcting for autocorrelation, we employ a solution proposed by Giordani and Halling (2019). The solution corrects the standard errors by multiplying them by $\sqrt{1+\frac{overlap}{2}}$. Then the adjusted t-statistics are derived in the following way:

$$t_{stat} = \frac{\beta}{SE} \tag{4.11}$$

$$t_{stat_{adj.}} = \frac{\beta}{SE * \sqrt{1+\frac{overlap}{2}}} = \frac{\beta}{SE} * \frac{1}{\sqrt{1+\frac{overlap}{2}}} = t_{stat} * \frac{1}{\sqrt{1+\frac{overlap}{2}}} \tag{4.12}$$

Where in our case, the overlap is equal to $overlap = holding\ period - 1$.

4.5 Model Selection Criterion

To compare the model fit, we employ an extended version of the Akaike Information Criterion (AIC) called Generalized Akaike Information Criterion (GAIC). The AIC penalises models for having more parameters, aiming to avoid overfitting.

The more parameters a model has, the better in-sample fit it can achieve. However, the overall variance of the model can suffer as a result, which implies that the out-of-sample performance of the model can become weaker. On the other hand, if we pick a rigid model, such as OLS, we introduce bias but reduce variance. To optimise the *bias-variance* trade-off, several information criteria have been developed, such as the AIC. One of the advantages of using AIC is that we can evaluate models of different dimensions without the need for out-of-sample tests, such as cross-validation. This is advantageous because we can train our model on the full sample size.

$$AIC = 2d - 2\ln(\hat{L}) \tag{4.13}$$

Where d is the number of estimated parameters in the model and $\hat{L}$ is the maximum value of the likelihood function for the model. The lower the AIC value, the better the model is considered to fit the data. Therefore, the model with the lowest AIC is preferred. To account for non-parametric models, GAMLSS uses GAIC as a default information criterion.

$$GAIC = k * df - 2\ln(\hat{L}) \qquad 4.14$$

Where k is a penalty term. Setting k = log (n), where n is the number of observations, yields the Bayesian Information Criterion (BIC). Setting $k = 2$ yields the regular AIC, the default GAMLSS setting, and the one we use to evaluate our models. Similar to AIC, the model with the lowest GAIC score is preferred.

Like the t-statistics mentioned in the previous section, the GAIC score will be off due to overlapping observations. Due to high autocorrelation, the log-likelihood will be overestimated, providing an artificially better model fit in terms of information criteria. When computing the information criteria, we use the same adjustment as model estimates, dividing the log-likelihood by the correction term.

$$GAIC_{adj.} = k * df - \frac{2\ln(\hat{L})}{\left(1 + \frac{overlap}{2}\right)} \qquad 4.15$$

Where in our case, the overlap is equal to $overlap = holding\ period - 1$.

4.6: Parameter Estimates and Risk Metrics

We estimate the distribution parameters $ST3 \sim (\mu, \sigma, \nu, \tau)$, $t \sim (\mu, \sigma, \nu)$, and $N \sim (\mu, \sigma)$ by maximum likelihood. The probability density functions of these distributions are available in closed form, and we can retrieve the PDF $\Pr[a \leq X \leq b] = \int_b^a f_x(x)dx$ based on the estimated parameters within the **gamlss** package. The PDF is used for simulation, from which the risk metrics can be computed empirically. The PDF can also be overlayed on the dataset's histogram to visualise the estimated distribution's fit.

Value-at-Risk and Expected Shortfall

We calculate the Value-at-Risk using the 'qSST()' function from the **gamlss** package, which computes the quantile function given specified distribution parameters. We are interested in the 1% Value-at-Risk, which means we compute the quantile of the loss distribution at the 1% level.

Lastly, to compute the Expected Shortfall, we use simulations. The expected shortfall is the expected loss given that Value-at-Risk is met or exceeded.

$$E[Loss|Loss \geq VaR_{\alpha}] \quad 4.16$$

Based on the distribution parameters, we simulate random draws from the SST distribution using the 'rSST()' function from the **gamlss** package. Next, we compute the average loss from these simulations but only consider the values equal to or exceeding the Value-at-Risk.

5. Results and Discussion

5.1: Model Parameters

In this section, we systematically evaluate the statistical models described in Chapter 4 using log covariates defined in Chapter 3. These models progressively increase in complexity by integrating more flexible distribution families with a greater number of distribution parameters. This procedure is similar to the one conducted by Giordani and Halling (2019) and allows us to evaluate the significance of the shape parameter.

	Normal			*Symmetric-T*			*Constant-Skew-T*			*Conditional-Skew-T*		
	LVL	**SLP**	**CRV**	**LVL**	**SLP**	**CRV**	**LVL**	**SLP**	**CRV**	**LVL**	**SLP**	**CRV**
$\beta_{0,\mu}$	0,06	0,06	0,06	0,07	0,07	0,07	0,14	0,14	0,15	0,15	0,15	0,15
adj. t-stat	*(3,68)*	*(3,75)*	*(3,69)*	*(4,66)*	*(4,34)*	*(4,32)*	*(5,41)*	*(5,30)*	*(5,77)*	*(5,76)*	*(5,78)*	*(5,76)*
$\beta_{1,\mu}$	-0,02	0,04	0,02	-0,02	0,03	0,01	-0,01	0,02	0,00			
adj. t-stat	*-(1,28)*	*(2,52)*	*(1,43)*	*-(1,43)*	*(2,19)*	*(0,79)*	*-(0,88)*	*(1,47)*	*(0,34)*			
$\beta_{0,\sigma}$	-1,81	-1,82	-1,81	-2,03	-1,99	-2,00	-2,13	-2,10	-2,12	-2,14	-2,13	-2,11
adj. t-stat	*-(27,98)*	*-(28,28)*	*-(28,01)*	*-(15,69)*	*-(16,64)*	*-(14,95)*	*-(15,78)*	*-(15,68)*	*-(15,27)*	*-(15,86)*	*-(16,50)*	*-(15,58)*
$\beta_{0,\nu}$							-0,40	-0,39	-0,42	-0,42	-0,41	-0,42
adj. t-stat							*-(2,92)*	*-(2,86)*	*-(3,12)*	*-(3,09)*	*-(3,05)*	*-(3,13)*
$\beta_{1,\nu}$										-0,07	0,14	0,08
adj. t-stat										*-(1,16)*	*(2,36)*	*(1,45)*
$\beta_{0,\tau}$				1,64	1,93	1,78	1,76	1,92	1,80	1,74	1,91	1,95
adj. t-stat				*(2,91)*	*(2,83)*	*(2,65)*	*(2,94)*	*(2,83)*	*(2,84)*	*(3,03)*	*(2,97)*	*(2,75)*
GAIC:	*- 86,91*	*- 91,46*	*- 87,29*	*- 90,04*	*- 92,76*	*- 88,62*	*- 98,34*	*- 99,72*	*- 97,66*	*- 98,92*	*- 103,12*	*- 99,58*

Table 5.1.1*: Illustrates parameter estimates for different statistical models with various distribution assumptions for the 12-month cumulative excess return horizon. The t-statistics provided in parenthesis are adjusted for overlapping observations, and so is the GAIC reported below each model. LVL, SLP, and CRV stand for level, slope, and curvature, respectively.*

	Normal			*Symmetric-T*			*Constant-Skew-T*			*Conditional-Skew-T*		
	LVL	**SLP**	**CRV**	**LVL**	**SLP**	**CRV**	**LVL**	**SLP**	**CRV**	**LVL**	**SLP**	**CRV**
$\beta_{0,\mu}$	0,03	0,03	0,03	0,03	0,03	0,04	0,06	0,07	0,07	0,07	0,07	0,07
adj. t-stat	*(3,57)*	*(3,59)*	*(3,55)*	*(5,01)*	*(4,90)*	*(4,94)*	*(4,73)*	*(4,86)*	*(5,18)*	*(5,15)*	*(5,22)*	*(5,19)*
$\beta_{1,\mu}$	-0,01	0,02	0,00	-0,02	0,02	0,00	-0,01	0,01	0,00			
adj. t-stat	*-(1,63)*	*(2,04)*	*(0,10)*	*-(2,35)*	*(2,26)*	*-(0,07)*	*-(1,72)*	*(1,93)*	*-(0,11)*			
$\beta_{0,\sigma}$	-2,15	-2,15	-2,14	-2,40	-2,37	-2,37	-2,41	-2,39	-2,39	-2,41	-2,41	-2,39
adj. t-stat	*-(47,16)*	*-(47,22)*	*-(47,04)*	*-(28,63)*	*-(29,84)*	*-(28,60)*	*-(28,35)*	*-(29,16)*	*-(28,24)*	*-(28,27)*	*-(28,86)*	*-(28,38)*
$\beta_{0,\nu}$							-0,22	-0,24	-0,25	-0,24	-0,24	-0,25
adj. t-stat							*-(2,45)*	*-(2,65)*	*-(2,88)*	*-(2,79)*	*-(2,81)*	*-(2,89)*
$\beta_{1,\nu}$										-0,05	0,08	0,01
adj. t-stat										*-(1,24)*	*(1,96)*	*(0,30)*
$\beta_{0,\tau}$				1,58	1,69	1,67	1,69	1,80	1,79	1,73	1,74	1,80
adj. t-stat				*(4,69)*	*(4,81)*	*(4,61)*	*(4,45)*	*(4,52)*	*(4,35)*	*(4,45)*	*(4,55)*	*(4,38)*
GAIC:	*- 345,31*	*- 346,77*	*- 342,68*	*- 359,87*	*- 359,61*	*- 354,55*	*- 364,07*	*- 364,81*	*- 361,18*	*- 362,75*	*- 365,10*	*- 361,26*

***Table 5.1.2**: Illustrates parameter estimates for different statistical models with various distribution assumptions for the 6-month cumulative excess return horizon. The t-statistics provided in parenthesis are adjusted for overlapping observations, and so is the GAIC reported below each model. LVL, SLP, and CRV stand for level, slope, and curvature, respectively.*

The benchmark model, Normal, is fitted using a normal distribution, where the mean is modelled as a function of individual yield curve factors, and the standard deviation is fixed. This model is equivalent to the OLS model in section 3.5, and the beta estimates from both models are identical. Looking at the slope factor, we can observe that a one standard deviation increase in *SLP* corresponds to an increase in 12-month (6-month) expected excess returns of 3.71% (1.53%).

Next, we expand our model using the Student's t-distribution, creating the Symmetric-T model. Since the Student's t-distribution is symmetric, the $\beta_{0,\mu}$ and $\beta_{1,\mu}$ estimates predict the mean of the return distribution. We can observe that the fit, measured by the adjusted GAIC, increases for all covariates compared to the Normal model. For the Symmetric-T model, the slope factor is still statistically significant in predicting the mode (mean). However, when allowing the residuals to follow fat-tailed distribution, the impact of a one standard deviation increase in *SLP* decreases to 3.25% for the 12-month time horizon but increases from 1.53% to 1.60% for the 6-month time horizon.

Constant-Skew-T extends the symmetric Student's t-distribution by including a skewness parameter. As indicated by the GAIC, allowing the distribution to capture asymmetry results in an improved fit for all covariates compared to the Symmetric-T model. The intercept $\beta_{0,\nu}$ implies negatively skewed distributions for all three covariates (-0.4 for *LVL*, -0.39 for *SLP*, and -0.42 for *CRV* for the 12-month time horizon). However, the mode estimate $\beta_{1,\mu}$ becomes less pronounced and more uncertain, as demonstrated by the decreasing t-statistic for both holding periods.

Lastly, the Conditional-Skew-T model introduces a time-varying shape estimate $\beta_{1,\nu}$ while keeping all other parameters constant and freely estimated. Compared to the Constant-Skew-T, this improves the model fit, as shown by decreasing values of GAIC for all covariates. Tables 5.1.1 and 5.1.2 show that *SLP* yields the highest estimates with the largest t-statistics for both time horizons. Conditioning *SLP*, we find positive shape estimates $\beta_{1,\nu}$ of 0,14 (0,08) for a 12-month holding period (6-month), with t-statistics above 1,96. A positive value for

$\beta_{1,\nu}$ indicates that as the *SLP* increases, the distribution becomes more positively skewed.

Motivated by the results of the slope covariate, we conduct a robustness check for different holding periods using the Conditional-Skew-T model. Table 5.1.3 presents $\beta_{1,\nu}$ estimates across holding periods between 1 and 24 months for all covariates. Examining the results, we observe that the slope covariate has the largest effect on the time-varying shape parameter $\beta_{1,\nu}$ compared to level and curvature covariates. The slope delivers both consistently higher $\beta_{1,\nu}$ estimates compared to other covariates, and the t-statistics are consistently above 1,96. Interestingly, for the slope covariate, the positive skewness increases with the holding period. As for the level and curvature factors, the estimates are small and not significant.

Conditional-Skew-T								
		Level	**Slope**	**Curvature**		**Level**	**Slope**	**Curvature**
$\beta_{1,\nu}$	**1M**	-0,04	0,06	0,01	**9M**	-0,06	0,11	0,06
t-stat		*-(2,20)*	*(3,18)*	*(0,62)*		*-(1,23)*	*(2,22)*	*(1,16)*
$\beta_{1,\nu}$	**3M**	-0,05	0,07	0,01	**12M**	-0,07	0,14	0,08
t-stat		*-(1,53)*	*(2,41)*	*(0,25)*		*-(1,16)*	*(2,36)*	*(1,45)*
$\beta_{1,\nu}$	**6M**	-0,05	0,08	0,01	**24M**	-0,12	0,17	0,05
t-stat		*-(1,24)*	*(1,96)*	*(0,30)*		*-(1,27)*	*(1,97)*	*(0,65)*

Table 5.1.3: *Provides beta estimates for the shape parameter,* $\beta_{1,\nu}$*, for various timeframes of cumulative returns, ranging from 1-month cumulative returns to 24-month cumulative returns.*

In summary, we can conclude from Table 5.1.1 and Table 5.1.2 that the Conditional-Skew-T yields the lowest GAIC score. This means that compared to other models, the model with the time-varying asymmetry best describes the cumulative log returns at the 12-month and 6-month horizon. This finding provides evidence of time-varying asymmetry in the returns. The main takeaway from the results in Tables 5.1.1 and 5.1.2 is that the distribution becomes more positively skewed when the slope of the yield curve is increasingly more upward sloping. This is captured by $\beta_{1,\nu} > 0$ for all time frames.

Furthermore, Table 5.1.3 shows that the slope factor provides consistent statistically significant estimates of the time-varying shape parameter $\beta_{1,\nu}$. In contrast, the other yield curve factors are typically not significant at a 90% confidence level, with only the level factor providing significant results at a 1-month holding period. Given the results in this section, we further conclude analysis only using the *SLP* covariate.

5.2: Predictive Distributions and Risk Return Dynamics

In the previous section, we examined the parameter estimates provided by our model and the implications of changes in yield curve factors on the return distribution. Here we present our principal findings of the thesis, particularly the predictive distributions of cumulative excess returns conditional on the slope and its implications for risk metrics. We explore the magnitude of the impact on the return distribution and the corresponding risk-return dynamics when high and low *SLP* readings are compared. Specifically, we focus on the predictive distribution conditional on *SLP* being two standard deviations from the mean. Tables 5.2.1, 5.2.4, 5.2.5, and 5.2.6 present the mean, mode, standard deviation, shape (only applies to Constant-Skew-T and Conditional-Skew-T), 1% Value-at-Risk (VaR), and 1% Expected Shortfall (ES) implied by the four models estimates presented in Section 5.1. Furthermore, Figure 5.2.2 and Figure 5.2.3 visualises the shifts in the distribution's conditional on the slope for the 12- and 6-month cumulative return period.

Predictive Distribution of Conditional-Skew-T

Conditional-Skew-T							
		Mean	***Mode***	***St.Dev***	***Shape***	***1% VaR***	***1% ES***
12M	**Low Values**	-1,47 %	14,63 %	19,70 %	0,50	-65,19 %	-84,22 %
	High Values	12,03 %	14,63 %	14,36 %	0,89	-27,05 %	-37,80 %
6M	**Low Values**	-0,23 %	6,87 %	12,73 %	0,66	-40,14 %	-52,91 %
	High Values	5,58 %	6,87 %	11,23 %	0,92	-24,76 %	-33,93 %

Table 5.2.1: *This table describes the key statistics from our conditional excess return distributions based on the Conditional-Skew-T model for the 12- and 6-month horizon. The table provides the mean, mode, standard deviation (St.Dev), the shape (skewness), the 1% Value-at-Risk, and the 1% Expected Shortfall. These statistics are reported for low values (2 standard deviations below the mean) and high values (2 standard deviations above the mean). All variables are reported in log.*

The mode, representing the most probable return of the distribution, is identical for high and low values as it is constructed to be fixed. However, if we compare it to the expected returns (mean), we observe that our model implies different measures when comparing low and high readings of *SLP*. While low *SLP* readings imply negative expected returns for both time frames (-1.47% for 12 months and -0.23% for 6 months), compared to positive expected returns when *SLP* is high (12.03% for 12 months and 5.58% for 6 months). The same observation can be made for the volatility (St.Dev), in that it decreases for high *SLP* (14.36% for 12 months and 11.23% for 6 months) compared to low *SLP* (19.70% for 12 months and 12.73% for 6 months). Both are characteristics favoured by investors. However, given how this model is constructed, these changes arise because of changes in the shape parameter.

As our primary interest lies in the shape of the distribution, further insights can be obtained by examining the shape metric, $0 < \nu_t < \infty$, which indicates the asymmetry of the distribution. This metric has an intuitive interpretation, in which $\nu = 1$ corresponds to a perfectly symmetric distribution, and a decrease in this value indicates increasing negative skewness. In Table 5.2.1, Figure 5.2.2, and Figure 5.2.3, we can observe that the distribution tends to become substantially more negatively skewed for low *SLP* readings. Interestingly, the return distribution nears symmetry when these values are high, as suggested by the shape parameter approaching the value of one and the mean converging with the mode. This means that the most probable outcome will be very close to the expected return.

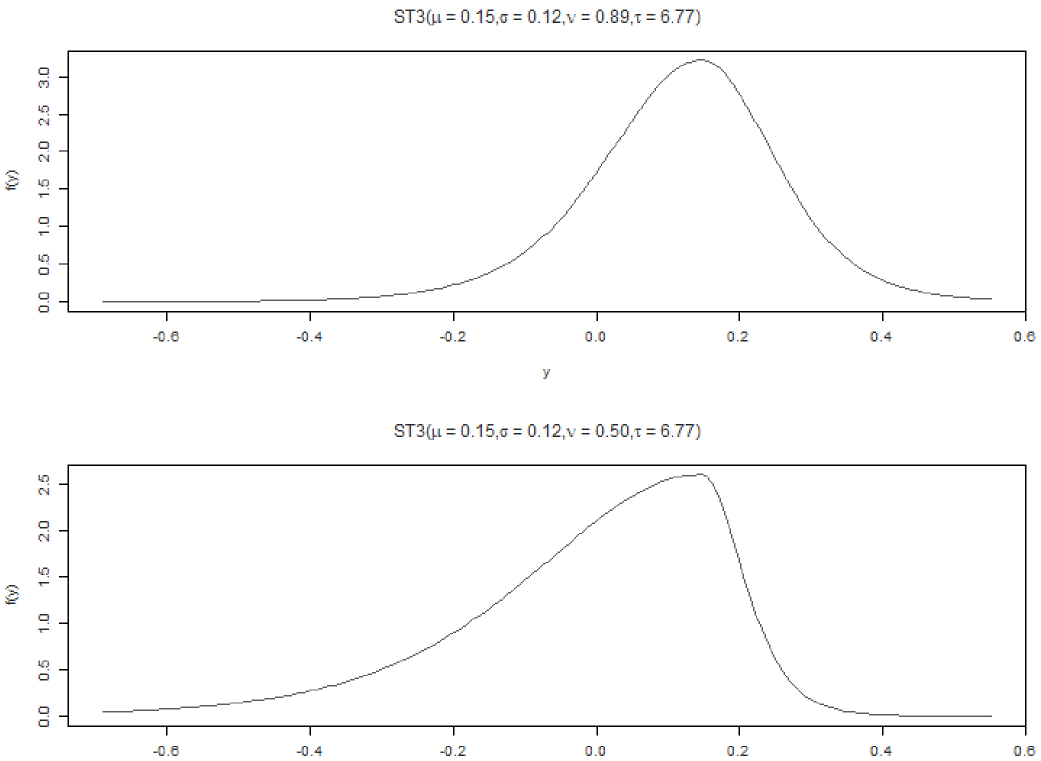

Figure 5.2.2*: 12M - The figure depicts the Conditional-Skew-T model's implied 12-month excess return conditional on SLP. Top graph is implied excess returns conditional on high SLP (2 standard deviations above the mean). Bottom graph is implied excess returns conditional on low SLP (2 standard deviations below the mean). Predictive distribution plots for the Normal, Symmetric-T, and Constant-Skew-T are provided in figures A.5.7 – A.5.9 in the appendix.*

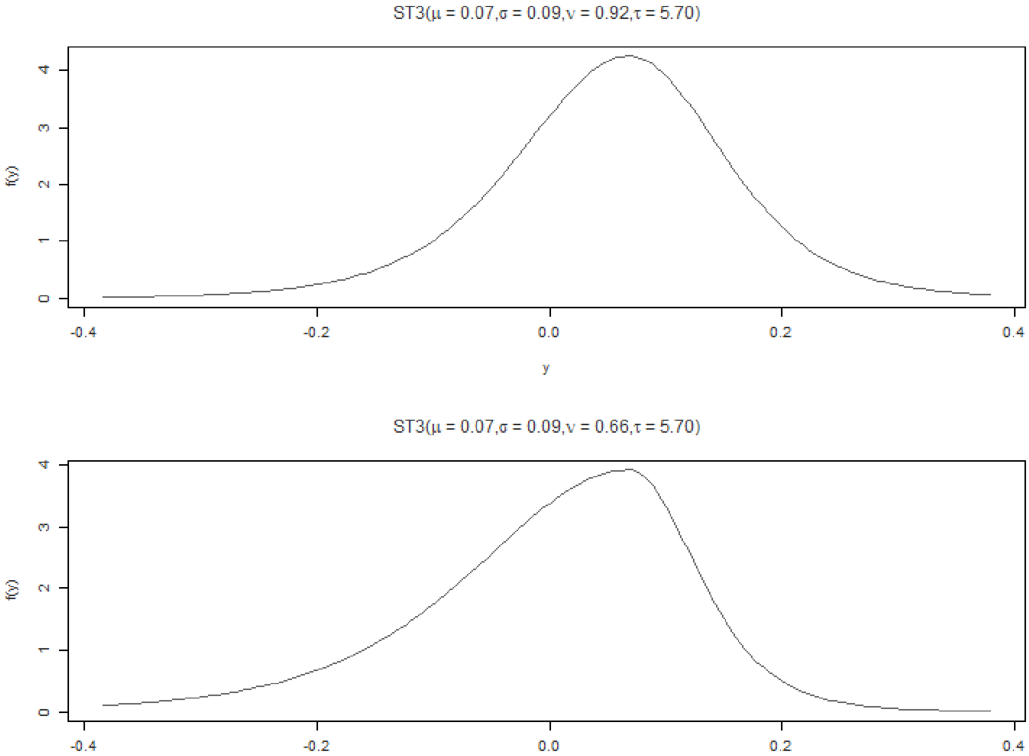

***Figure 5.2.3**: 6M - The figure depicts the Conditional-Skew-T model's implied 6-month excess return conditional on SLP. Top graph is implied excess returns conditional on high SLP (2 standard deviations above the mean). Bottom graph is implied excess returns conditional on low SLP (2 standard deviations below the mean). Predictive distribution plots for the Normal, Symmetric-T, and Constant-Skew-T are provided in figures A.5.10 – A.5.12 in the appendix.*

Risk Return Dynamics and Comparison with the Normal Distribution

Accurately defining the return distribution is particularly important when dealing with risk metrics. To illustrate the significance of this point, we compare the predictive distributions derived from the Conditional-Skew-T and the normal distribution. Two key factors drive our decision to compare the Conditional-Skew-T with the normal distribution: firstly, it allows us to underscore the considerable underestimations in risk metrics that arise from misspecification of the distribution; secondly, it highlights a common pitfall in predictive modelling – namely, the unwarranted reliance on the normal distribution for financial time series data which in fact, does not conform to a normal.

Normal						
		Mean	***Mode***	***St.Dev***	***1% VaR***	***1% ES***
12M	**Low Values**	-1,89 %	-1,89 %	16 %	-39,40 %	-44,84 %
	High Values	12,97 %	12,97 %	16 %	-24,54 %	-29,94 %
6M	**Low Values**	-0,36 %	-0,36 %	12 %	-27,45 %	-31,38 %
	High Values	5,73 %	5,73 %	12 %	-21,36 %	-25,26 %

Table 5.2.4: *This table describes the key statistics from our conditional excess return distributions based on the Normal model for the 12- and 6-month horizon. The table provides the mean, mode, standard deviation (St.Dev), the 1% Value-at-Risk, and the 1% Expected Shortfall. These statistics are reported for low values (2 standard deviations below the mean) and high values (2 standard deviations above the mean). All variables are reported in log.*

Risk metrics derived from a symmetric, normal distribution (such as standard deviation, 1% VaR, 1% ES) are significantly underestimated compared to Conditional-Skew-T metrics that account for larger-than-normal tails and asymmetry. For instance, considering the 1% VaR for the 12-month cumulative returns when *SLP* is low, the Normal model projects a 1% probability of a -39.40% loss (-32.57% in arithmetic returns). However, with a more flexible distributional assumption, this potential loss surges to -65.19% (-47.89% in arithmetic returns). The discrepancy becomes even more pronounced when considering the 1% ES.

Risk-Return Dynamics and t-Distributions

We compare the Symmetric-Skew-T and Constant-Skew-T with the Conditional-Skew-T model to address the nuances in risk-return dynamics. Across all models, we observe that expected returns are considerably higher when *SLP* is high than when it is low. Also, the expected returns are similar across all models, including the Normal model.

From Tables 5.2.5 and 5.2.6 we observe that for high *SLP* readings the 1% ES is -37.80% for the Conditional-Skew-T model, compared to -54.95% for the Constant-Skew-T model. For low slope readings, the 1% ES is at -84.22% for the Conditional-Skew-T model, compared to -62.84% for the Constant-Skew-T model.

These results align with the findings of Giordani & Halling (2019), that used a similar approach when predicting return distributions conditional on the CAPE valuation metric. Specifically, for higher slope values, both the Symmetric-Skew-T and Constant-Skew-T models tend to overestimate risk metrics such as standard deviation, VaR, and ES while underestimating them for lower slope readings. These findings highlight that selecting the appropriate distributional assumption is key to accurately capturing financial markets' risk and return dynamics and has economically significant implications for risk management.

Symmetric-T						
		Mean	*Mode*	*St.Dev*	*1% VaR*	*1% ES*
12M	Low Values	0,19 %	0,19 %	16 %	-41,02 %	-51,37 %
	High Values	13,23 %	13,23 %	16 %	-27,98 %	-38,66 %
6M	Low Values	0,26 %	0,26 %	12 %	-30,09 %	-38,96 %
	High Values	6,64 %	6,64 %	12 %	-23,72 %	-32,92 %

Table 5.2.5: *This table describes the key statistics from our conditional excess return distributions based on the Symmetric-T model for the 12- and 6-month horizon. The table provides the mean, mode, standard deviation (St.Dev), the shape (skewness), the 1% Value-at-Risk, and the 1% Expected Shortfall. These statistics are reported for low values (2 standard deviations below the mean), and high values (2 standard deviations above the mean). All variables are reported in log.*

Constant-Skew-T							
		Mean	*Mode*	*St.Dev*	*Shape*	*1% VaR*	*1% ES*
12M	Low Values	1,11 %	9,99 %	16 %	0,68	-48,65 %	-62,84 %
	High Values	9,56 %	18,34 %	16 %	0,68	-40,30 %	-54,95 %
6M	Low Values	-0,14 %	3,95 %	12 %	0,79	-34,19 %	-44,24 %
	High Values	5,49 %	9,51 %	12 %	0,79	-28,63 %	-39,03 %

Table 5.2.6: *This table describes the key statistics from our conditional excess return distributions based on the Constant-Skew-T model for the 12- and 6-month horizon. The table provides the mean, mode, standard deviation (St.Dev), the shape (skewness), the 1% Value-at-Risk, and the 1% Expected Shortfall. These statistics are reported for low values (2 standard deviations below the mean), and high values (2 standard deviations above the mean). All variables are reported in log.*

5.3: Predictive Distributions for Pre- and Post-1985

In this section, we focus on the recent findings of Estrella et al. (2003), which suggest that the predictive power of the slope for output growth in the US has diminished after the mid-1980s. However, in our thesis, we aim to investigate whether the predictive power is also diminished for the shape of the return distribution in US equity markets. As concluded in the previous sections, we provide evidence that the Conditional-Skew-T model best describes the returns in our dataset. Additionally, among the yield curve factors examined, the slope factor exhibited the highest level of robustness.

To test the hypothesis regarding the diminished predictive power of the slope, we divided the sample into pre-1985 and post-1985 subsamples, as outlined in Chapter 3.3. For each sub-sample, we estimate the Conditional-Skew-T model and run the predictive distribution on the estimated parameters to grasp any changes in the risk-return characteristics between the two periods. However, since the pre-1985 sample is about half the sample size of the post-1985 sample, the estimates are less reliable for the former. Additionally, the variance in our explanatory variable is smaller in the post-1985 sample (1.2806) compared to the pre-1985 sample (1.5726), which also reduces statistical precision.

	12 months		*6 months*	
Conditional-Skew-T	**Pre-1985**	**Post-1985**	**Pre-1985**	**Post-1985**
$\beta_{0,\mu}$	0,10	0,17	0,02	0,08
adj. t-stat	*(4,56)*	*(6,04)*	*(0,90)*	*(5,43)*
$\beta_{0,\sigma}$	-1,93	-2,26	-2,21	-2,55
adj. t-stat	*-(17,54)*	*-(12,95)*	*-(17,49)*	*-(23,20)*
$\beta_{0,\nu}$	-0,29	-0,45	-0,07	-0,29
adj. t-stat	*-(3,03)*	*-(2,62)*	*-(0,48)*	*-(2,52)*
$\beta_{1,\nu}$	0,18	0,08	0,14	0,02
adj. t-stat	*(2,36)*	*(1,02)*	*(2,21)*	*(0,31)*
$\beta_{0,\tau}$	58,50	1,44	3,02	1,43
adj. t-stat	*(38665710,99)*	*(2,55)*	*(1,50)*	*(3,78)*

Table 5.3.1: *Illustrates parameter estimates for the Conditional-Skew-T for 6-month and 12-month cumulative excess return horizon when estimated on pre-1985 and post-1985 sub-samples. The t-statistics provided in parenthesis are adjusted for overlapping observations*.

Table 5.3.1 shows that the $\beta_{1,\nu}$ coefficients become much smaller in the post-1985 sample compared to the pre-1985 sample for both holding periods. While the $\beta_{1,\nu}$ estimate almost halves for the 12-month holding period (0,18 to 0,08), the change is more pronounced for the 6-month holding period, where the estimate decrease from 0,14 to 0,02. In addition, the t-statistics for both holding periods fall below the 90% significance level in the post-1985 sample. This provides evidence, in line with related literature, that the predictive power of the slope as a predictor for stock returns has diminished since mid-1985.

The predictive distribution of the pre-1985 sample illustrated in Table 5.3.2 shows large discrepancies in the risk-return characteristics between high and low slope values. The mean changes from negative (-6,56% for 12-month and -4,23% for 6-month) to positive (11,11% for 12-month and 6,32% for 6-month) for both holding periods. Similarly, the standard deviation decreases for high *SLP* readings for both holding periods, signalling more favourable risk-return characteristics for investors. Interestingly, the skewness of the distribution turns positive for high slope values with the shape parameter $\nu > 1$. As expected, the VaR and ES values become much smaller for high slope values than under conditions with low or negative term spreads

Conditional-Skew-T							
Pre-1985		*Mean*	*Mode*	*St.Dev*	*Shape*	*1% VaR*	*1% ES*
12M	Low Values	-6,56 %	9,53 %	18,82 %	0,52	-59,21 %	-67,98 %
	High Values	11,11 %	9,53 %	14,54 %	1,07	-21,55 %	-26,15 %
6M	Low Values	-4,23 %	2,37 %	12,62 %	0,70	-39,05 %	-45,78 %
	High Values	6,32 %	2,37 %	11,98 %	1,24	-18,99 %	-22,97 %

Table 5.3.2: *This table describes the key statistics from our conditional excess return distributions based on the Conditional-Skew-T model for the 12- and 6-month horizon for the pre-1985 sub-sample. The table provides the mean, mode, standard deviation (St.Dev), the 1% Value-at-Risk, and the 1% Expected Shortfall. These statistics are reported for low values (2 standard deviations below the mean) and high values (2 standard deviations above the mean). All variables are reported in log.*

A drastic change in the predictive distribution can be observed in the post-1985 sub-sample in Table 5.3.3. Though the discrepancies are still economically significant for the 12-month holding period, the 6-month holding period displays the minimal change between high and low slope readings. While the shape parameter changes from 0,54 to 0,75 for the 12-month holding period, it remains almost unchanged for the 6-month holding period. Although *SLP* values indicate larger discrepancies for the 12-month holding period, the slope no longer seems to predict stock market contraction (negative cumulative expected returns) as it did for the full sample in Table 5.2.1 and the pre-1985 sample in Table 5.3.2.

Conditional-Skew-T							
Post-1985		*Mean*	*Mode*	*St.Dev*	*Shape*	*1% VaR*	*1% ES*
12M	Low Values	3,18 %	16,76 %	19,34 %	0,54	-63,16 %	-90,64 %
	High Values	10,88 %	16,76 %	15,55 %	0,75	-37,39 %	-57,21 %
6M	Low Values	3,09 %	8,32 %	11,87 %	0,72	-34,78 %	-49,85 %
	High Values	4,28 %	8,32 %	11,59 %	0,77	-31,31 %	-45,90 %

Table 5.3.3: *This table describes the key statistics from our conditional excess return distributions based on the Conditional-Skew-T model for the 12- and 6-month horizon for the post-1985 sub-sample. The table provides the mean, mode, standard deviation (St.Dev), the 1% Value-at-Risk, and the 1% Expected Shortfall. These statistics are reported for low values (2 standard deviations below the mean) and high values (2 standard deviations above the mean). All variables are reported in log.*

5.4: International Evidence

As additional evidence to ensure the robustness of our findings, we conduct the analysis on international data. Specifically, we use data from the UK, France, Australia, and Germany. For each country, we construct the slope using the log yield of the 10-year and 2-year government bonds.

$$SLP_{(INT)} = \ln(y_{10yr} - y_{2yr}) \qquad 5.1$$

The slope construction deviates from the main analysis due to data availability in international markets. To control for this change, we also include the US market, with the slope constructed the same way as other international markets. For equity returns, we use the total return index in each country, commencing 29th of May 1992, to have comparable lengths for all markets. The appendix in Tables A.5.1 – A.5.6 includes a complete overview of model estimates.

	Conditional-Skew-T	US	UK	DE	AUS	FR
12M	$\beta_{1,\nu}$	0,08	0,20	0,27	0,16	0,20
	adj. t-stat	*(0,81)*	*(2,33)*	*(2,86)*	*(2,00)*	*(2,48)*
6M	$\beta_{1,\nu}$	0,05	0,12	0,23	0,09	0,19
	adj. t-stat	*(0,80)*	*(1,91)*	*(3,43)*	*(1,51)*	*(2,98)*

Table 5.4.1: *This table describes the shape parameter $\beta_{1,\nu}$ and its t-statistics estimated by the Conditional-Skew-T. The t-statistics are adjusted for overlapping observations. The shape parameter is estimated on total return indexes in different international markets for 12-month and 6-month holding periods.*

Our analysis in Chapter 5.3 reveals evidence of limited predictive performance for the slope in US markets for the post-1985 subsample. This is in line with the findings in Table 5.4.1, namely that the slope has limited predictive accuracy for the post-1992 sample. However, the results suggest that the slope remains relevant for predicting the distribution shape in other developed markets post-1992.

Furthermore, $\beta_{1,\nu}$ estimates are always positive. This confirms that the slope of the yield curve and the skewness of cumulative equity returns are positively related. Interestingly, all markets in the sample show a much higher slope estimate than the US market, with the German market displaying the most pronounced effect of 0,27 for a 12-month holding period, compared to only 0,08 for the US market. The t-statistics in non-US markets are all above 2 for the 12-month holding period, compared to only 0,81 for the US market, displaying a higher level of coefficient estimate certainty in non-US developed markets.

Table 5.4.2. and Table 5.4.3 confirms that the Conditional-Skew-T model is supported in the data for all non-US markets in the sample. This is illustrated by the Conditional-Skew-T model having a lower GAIC score compared to the Constant-Skew-T and Symmetric-T models. However, for the US, the Constant-Skew-T provides a lower GAIC than the Conditional-Skew-T for the 6-month holding period.

GAIC 12M	US	UK	DE	AUS	FR
Conditional-Skew-T	*-43,52*	*- 63,52*	*- 16,54*	*- 64,53*	*- 20,82*
Constant-Skew-T	*-42,91*	*- 60,52*	*- 13,68*	*- 61,41*	*- 17,98*
Symmetric-T	*-39,33*	*- 57,53*	*- 11,84*	*- 60,75*	*- 17,06*

Table 5.4.2: *This table lists the GAIC values for the Conditional-Skew-T, Constant-Skew-T and Symmetric-T models. Models are estimated on total return indexes in different international markets for the 12-month holding period. The GAIC is adjusted for overlapping observations.*

GAIC 6M	US	UK	DE	AUS	FR
Conditional-Skew-T	*-193,20*	*- 223,07*	*- 121,85*	*- 219,29*	*- 144,21*
Constant-Skew-T	*-193,32*	*- 221,29*	*- 115,57*	*- 217,20*	*- 138,89*
Symmetric-T	*-189,95*	*- 217,77*	*- 113,74*	*- 214,74*	*- 136,06*

Table 5.4.3: *This table lists the GAIC values for the Conditional-Skew-T, Constant-Skew-T and Symmetric-T models. Models are estimated on total return indexes in different international markets for the 6-month holding period. The GAIC is adjusted for overlapping observations.*

In summary, we discovered evidence from the sample of international developed markets that contradicts the primary findings for the US market in the post-1985 sample. The slope covariate demonstrates stronger in-sample predictive power for the shape of equity returns in international markets, with a coefficient of at least 0,16 and a t-statistics equal to or above 2.0 for the 12-month holding period. Additionally, as documented by the GAIC score, our findings suggests that the Conditional-Skew-T model, with its time-varying shape, provides a better fit to the data for non-US markets compared to the other models.

5.5: Limitations and Assumptions

It is important to acknowledge the limitations and assumptions inherent in the analysis. Firstly, while there may be other variables, or combinations of variables, that could offer a better fit for the data, this study focuses on understanding how the return distribution changes based on level, slope, and curvature readings individually. Therefore, the decision to include these variables in the analysis is driven by the research objectives and not by the goal of constructing the best-fitting model.

Lastly, it is worth noting that the entire analysis is conducted solely on in-sample data. This means the findings and conclusions are based solely on the observed data used for analysis. Extrapolating these results to out-of-sample or future data should be done with caution, as the model’s performance and assumptions might be less effective when applied to new or unseen data.

Including out-of-sample testing and validation would provide a more robust assessment of the model's predictive capabilities. However, since the objective is not to construct a model for the best fit but to explore the general dynamics of the return distribution in relation to changing yield curve factors, the emphasis on out-of-sample analysis is relatively less important. While out-of-sample analysis would be crucial for accurate forecasting, in this context, the focus is on understanding the broader patterns and relationships between the yield curve factors and the changing dynamics of the return distribution.

6. Conclusion

In this paper, we investigated yield curve factors' efficacy in predicting the distribution of cumulative 12-month and 6-month total log returns in the aggregate US market and other developed markets. Our key results indicate that the slope of the yield curve is the most influential factor affecting the shape of the return distribution, compared to other factors such as level and curvature. Specifically, as the slope becomes increasingly upward-sloping, the return distribution approaches symmetry, while a lower slope leads to the distribution becoming more negatively skewed.

However, as depicted in the literature, the significance of using the slope to forecast output growth has diminished after the mid-1980s, and we find a similar reduction in the slope's efficacy in forecasting the shape of US equity return distribution. Despite the reduced effect of the slope in US markets post-mid-1980s, we find evidence that the slope is still useful for predicting the shape of the distribution in some other developed markets. Specifically in the UK, Germany, France, and Australia, the effect seems to persist.

While existing literature emphasises the prediction of expected returns, we focus on the importance of accurately describing the shape of the distribution for risk management purposes. To highlight this importance, we derive risk metrics, such as Value-at-Risk and Expected Shortfall and compare them across several models with different distributional assumptions. Compared to a model which allows for time-varying asymmetry, we find that using a normal distribution, a symmetric distribution, or one with constant skewness drastically underestimated the risk metrics when the yield curve is flat or downward sloping and overestimated them when it was upward sloping.

By moving beyond traditional mean-based models, such as the OLS, we provided a more comprehensive understanding of how changes in yield curve factors influence the shape of the distribution of cumulative equity excess returns. By incorporating GAMLSS, a novel framework for modelling non-normality and asymmetry, we contribute to existing literature and offer a fresh perspective on flexible distribution modelling of financial return series.

7. References

Albuquerque, R. (2012). *Skewness in Stock Returns: Reconciling the Evidence on Firm Versus Aggregate Returns*. The Review of Financial Studies, 25(5), 1630–1673. http://www.jstor.org/stable/41485537

Ang A. (2014). *Asset management a systematic approach to factor investing*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199959327.001.0001

Ang, A., Piazzesi, M., Wei, M. (2006). *What does the yield curve tell us about GDP growth?*. Journal of Econometrics, 131(1–2), 359-403. https://doi.org/10.1016/j.jeconom.2005.01.032

Bernanke, B. S., & Gertler, M. (1995). Inside the Black Box: *The Credit Channel of Monetary Policy Transmission*. The Journal of Economic Perspectives, 9(4), 27–48. http://www.jstor.org/stable/2138389

Brooks, C. (2019). *Introductory Econometrics for Finance* (4th ed.). Cambridge: Cambridge University Press. doi:10.1017/9781108524872

Brunnermeier M. K., & Pedersen, L. H. (2009). *Market Liquidity and Funding Liquidity*. The Review of Financial Studies, 22(6), 2201–2238. https://doi.org/10.1093/rfs/hhn098

Campbell, J.Y. (1987). *Stock returns and the term structure*. Journal of Financial Economics, 18(2), 373-399. https://doi.org/10.1016/0304-405X(87)90045-6

Cenesizoglu, T., & Timmermann, A. (2008). *Is the Distribution of Stock Returns Predictable?* Available at SSRN: https://ssrn.com/abstract=1107185

Chava, S., Purnanandam, A. (2011). *The effect of banking crisis on bank-dependent borrowers*. Journal of Financial Economics, 99(1), 116-135. https://doi.org/10.1016/j.jfineco.2010.08.006

Chava, S.,Gallmeyer, M. F., Park, H. (2015). *Credit Conditions and Stock Return Predictability*. Journal of Monetary Economics, 74, 117-132. http://dx.doi.org/10.2139/ssrn.1571958

Chen, N.-F. (1991). *Financial Investment Opportunities and the Macroeconomy*. The Journal of Finance, *46*(2), 529–554. https://doi.org/10.2307/2328835

Cochrane, J. H., Piazzesi., M. (2005). *Bond Risk Premia*. American Economic Review, 95(1), 138-160. https://www.aeaweb.org/articles?id=10.1257/0002828053828581

D'Agostino, A., Giannone, D., and Surico, P. (2006). *(Un)Predictability and Macroeconomic Stability*. ECB Working Paper No. 605. http://dx.doi.org/10.2139/ssrn.890990

Dotsey, M. (1998). *The Predictive Content of the Interest Rate Term Spread for Future Economic Growth*. FRB Richmond Economic Quarterly, 84(3), 31-51. https://ssrn.com/abstract=2126284

Estrella, A. (2005). *Why Does the Yield Curve Predict Output and Inflation?*. The Economic Journal, 115(505), 722-744. https://doi.org/10.1111/j.1468-0297.2005.01017.x

Estrella, A., Rodrigues, A. R., & Schich, S. (2003). *How Stable Is the Predictive Power of the Yield Curve? Evidence from Germany and the United States*. The Review of Economics and Statistics, 85(3), 629–644. http://www.jstor.org/stable/3211702

Fama, E. F. (1965). *The Behavior of Stock-Market Prices*. The Journal of Business, 38(1), 34–105. http://www.jstor.org/stable/2350752

Fama, E. F., & French, K. R. (1989). *Business conditions and expected returns on stocks and bonds*. Journal of Financial Economics, 25(1), 23-49. https://doi.org/10.1016/0304-405X(89)90095-0

Fernandez, C., & Mark F. J. Steel. (1998). *On Bayesian Modeling of Fat Tails and Skewness*. Journal of the American Statistical Association, 93(441), 359–371. https://doi.org/10.2307/2669632

Giordani, P., & Halling, M. (2019). *Valuation Ratios and Shape Predictability in the Distribution of Stock Returns*. Swedish House of Finance Research Paper No. 17-5. https://ssrn.com/abstract=2887156

Hastie, T., & Tibshirani, R. (1986). *Generalized Additive Models*. Statist. Sci. 1 (3) 297 – 310. https://doi.org/10.1214/ss/1177013604

Ilmanen A. (2011). *Expected returns: an investor's guide to harvesting market rewards* (1st ed.). John Wiley.

Mandelbrot, B. (1963). *The Variation of Certain Speculative Prices*. The Journal of Business, 36(4), 394–419. http://www.jstor.org/stable/2350970

Mandelbrot, B., & Taylor, H. M. (1967). *On the Distribution of Stock Price Differences*. Operations Research, 15(6), 1057–1062. http://www.jstor.org/stable/168611

Nelder, J. A., & Wedderburn, R. W. M. (1972). Generalized Linear Models. Journal of the Royal Statistical Society, Series A (General), 135(3), 370–384. https://doi.org/10.2307/2344614

Plosser, C. I., Rouwenhorst, K., G. (1994). *International term structures and real economic growth*. Journal of Monetary Economics, 33(1), 133-155. https://doi.org/10.1016/0304-3932(94)90017-5

Rigby, R. A., & Stasinopoulos, D. M. (2005). *Generalized Additive Models for Location, Scale and Shape*. Journal of the Royal Statistical Society. Series C (Applied Statistics), 54(3), 507–554. http://www.jstor.org/stable/3592732

Rigby, R.A., Stasinopoulos, M.D., Heller, G.Z., Voudouris, V., & De Bastiani, F. (2017). *Flexible Regression and Smoothing: Using GAMLSS in R* (1st ed.). Chapman and Hall/CRC. https://doi.org/10.1201/b21973

Schwert, G. W. (1990). *Stock Returns and Real Activity: A Century of Evidence*. The Journal of Finance, 45(4), 1237–1257. https://doi.org/10.2307/2328722

Veronesi, P. (2010). *Fixed Income Securities: Valuation, Risk, and Risk Management* (1st ed.). John Wiley.

Welch, I., & Goyal, A. (2008). *A Comprehensive Look at the Empirical Performance of Equity Premium Prediction*. The Review of Financial Studies, 21(4), 1455–1508. http://www.jstor.org/stable/40056859

8. Appendix

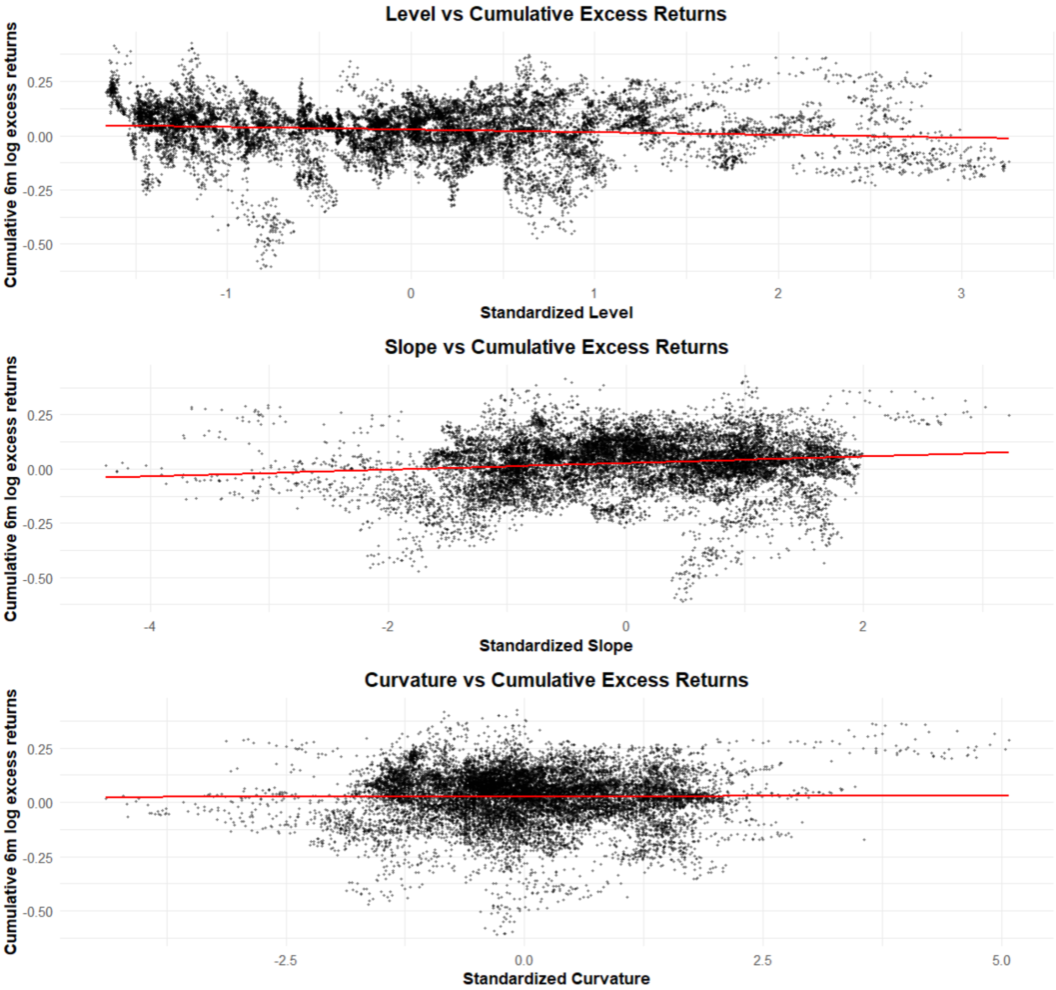

Figure A.3.1: *Scatter plots of standardised yield curve factors (x-axis) vs cumulative 6-month excess log returns (y-axis) with a fitted simple linear regression line. The top plot shows the standardised level vs returns, the middle plot shows the standardised slope vs returns, and the bottom plot shows the standardised curvature vs returns.*

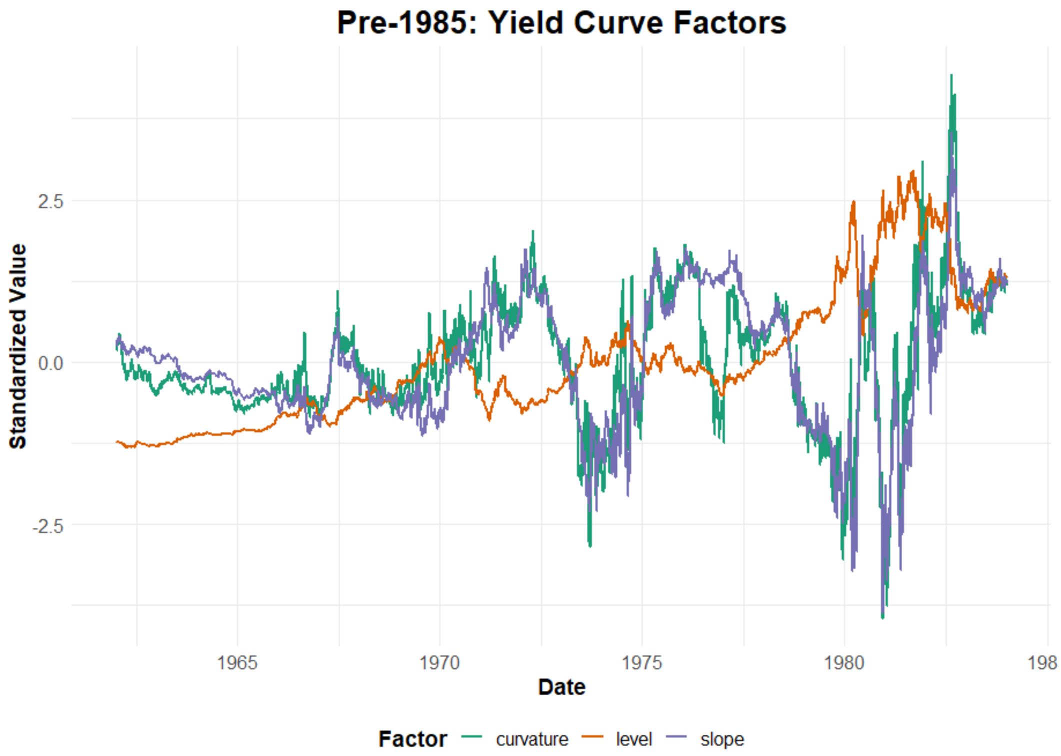

Figure A.3.2: *Time series of standardised values of yield curve factors. A pre-1985 sample is used. The orange line depicts the standardised level factor, the purple line shows the standardised slope factor, and the green line highlights the standardised curvature factor.*

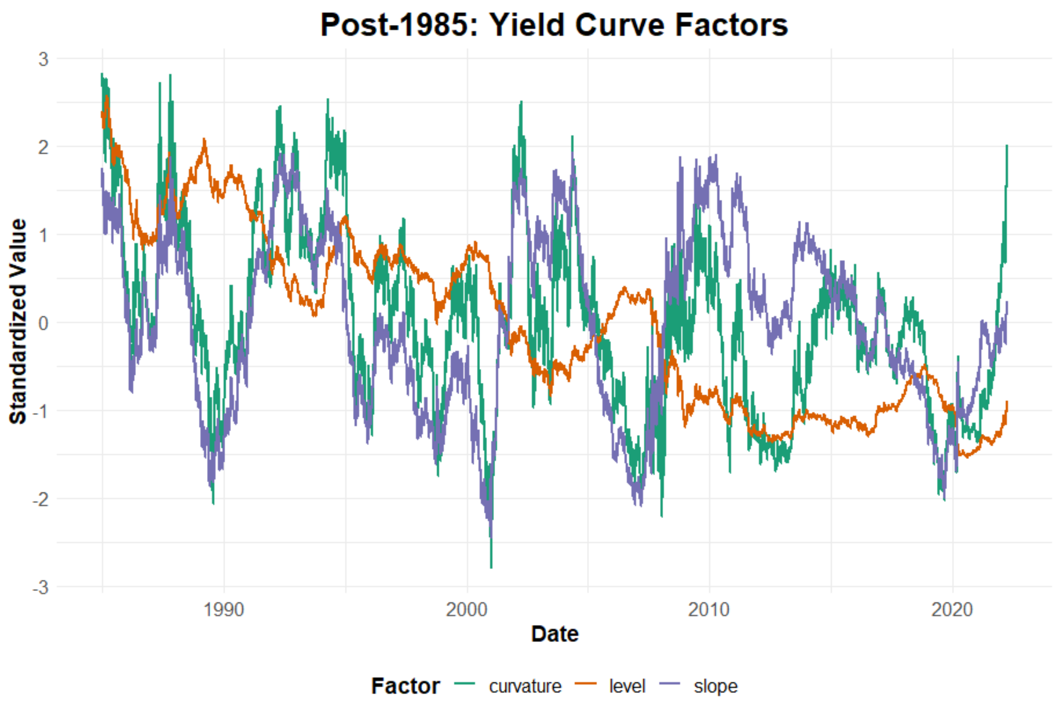

Figure A.3.3: *Time series of standardised values of yield curve factors. A post-1985 is used. The orange line depicts the standardised level factor, the purple line shows the standardised slope factor, and the green line highlights the standardised curvature factor.*

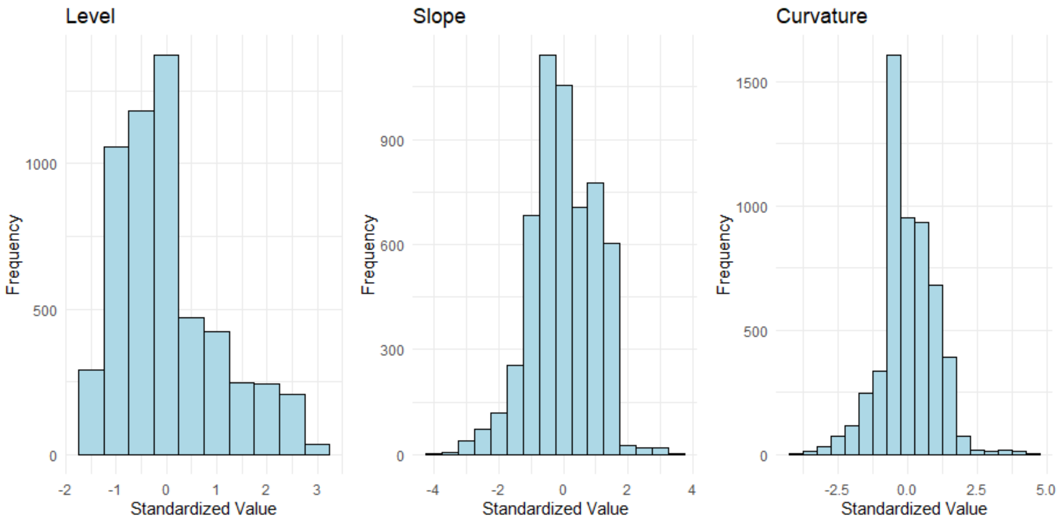

Figure A.3.4: *This figure presents histograms of the standardised level, slope, and curvature factors for the pre-1985 sample. Each histogram visually illustrates the distribution of its respective yield curve factor.*

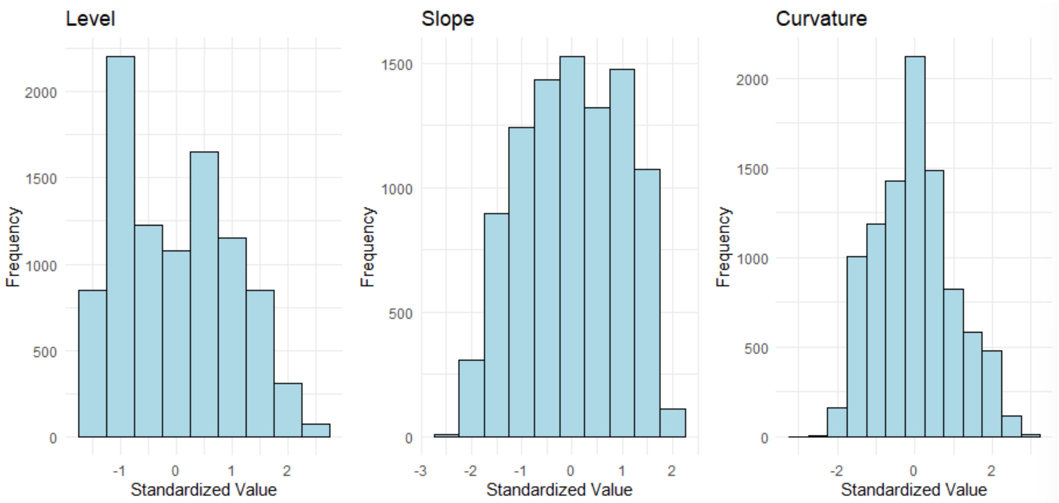

Figure A.3.5: *This figure presents histograms of the standardised level, slope, and curvature factors for the post-1985 sample. Each histogram visually illustrates the distribution of its respective yield curve factor.*

Conditional-Skew-T (12M)	US	UK	DE	AUS	FR
$\beta_{0,\mu}$	0.1653	0.1503	0.2152	0.1707	0.1994
adj. t-stat	*(5.29)*	*(4.72)*	*(5.30)*	*(4.32)*	*(3.92)*
$\beta_{0,\sigma}$	-2.2288	-2.1851	-1.9162	-2.1658	-1.7810
adj. t-stat	*-(11.45)*	*-(12.35)*	*-(10.45)*	*-(13.37)*	*-(14.86)*
$\beta_{0,\nu}$	-0.4169	-0.4081	-0.5000	-0.3928	-0.4266
adj. t-stat	*-(2.28)*	*-(2.28)*	*-(2.81)*	*-(1.78)*	*-(2.25)*
$\beta_{1,\nu}$	0.0772	0.2007	0.2718	0.1580	0.1993
adj. t-stat	*(0.81)*	*(2.33)*	*(2.86)*	*(2.00)*	*(2.48)*
$\beta_{0,\tau}$	1.3471	2.8915	2.1004	3.6076	11.8271
adj. t-stat	*(2.24)*	*(1.19)*	*(1.94)*	*(0.76)*	*(0.04)*

Table A.5.1: *Illustrates parameter estimates for various international markets using the Conditional-Skew-T model for the 12-month holding period. The t-statistics provided in parenthesis are adjusted for overlapping observations.*

Conditional-Skew-T (6M)	US	UK	DE	AUS	FR
$\beta_{0,\mu}$	0.0821	0.0759	0.1086	0.0863	0.0999
adj. t-stat	*(5.18)*	*(5.24)*	*(5.04)*	*(5.21)*	*(5.23)*
$\beta_{0,\sigma}$	-2.5858	-2.6405	-2.1939	-2.5820	-2.2778
adj. t-stat	*-(20.87)*	*-(21.01)*	*-(18.33)*	*-(21.58)*	*-(18.19)*
$\beta_{0,\nu}$	-0.2892	-0.2827	-0.3312	-0.2842	-0.3190
adj. t-stat	*-(2.30)*	*-(2.30)*	*-(2.72)*	*-(2.16)*	*-(2.70)*
$\beta_{1,\nu}$	0.0505	0.1165	0.2269	0.0865	0.1864
adj. t-stat	*(0.80)*	*(1.91)*	*(3.43)*	*(1.51)*	*(2.98)*
$\beta_{0,\tau}$	1.3172	1.5687	1.9250	1.8885	1.8944
adj. t-stat	*(3.46)*	*(3.18)*	*(3.05)*	*(2.85)*	*(2.97)*

Table A.5.2: *Illustrates parameter estimates for various international markets using the Conditional-Skew-T model for the 6-month holding period. The t-statistics provided in parenthesis are adjusted for overlapping observations.*

Constant-Skew-T (12M)	US	UK	DE	AUS	FR
$\beta_{0,\mu}$	0.1655	0.1507	0.1927	0.1586	0.1710
adj. t-stat	*(5.19)*	*(4.21)*	*(3.54)*	*(4.40)*	*(2.89)*
$\beta_{1,\mu}$	0.0035	0.0256	0.0549	-0.0066	0.0430
adj. t-stat	*(0.20)*	*(1.38)*	*(2.34)*	*-(0.31)*	*(1.76)*
$\beta_{0,\sigma}$	-2.2267	-2.1758	-1.8059	-2.2898	-1.7086
adj. t-stat	*-(11.02)*	*-(11.68)*	*-(10.42)*	*-(11.68)*	*-(14.99)*
$\beta_{0,\nu}$	-0.4200	-0.4104	-0.4087	-0.3327	-0.3381
adj. t-stat	*-(2.26)*	*-(2.10)*	*-(1.97)*	*-(1.56)*	*-(1.66)*
$\beta_{0,\tau}$	1.3317	2.4450	2.2574	1.6065	12.4048
adj. t-stat	*(2.17)*	*(1.50)*	*(1.77)*	*(2.06)*	*(0.02)*

Table A.5.3: *Illustrates parameter estimates for various international markets using the Constant-Skew-T model for the 12-month holding period. The t-statistics provided in parenthesis are adjusted for overlapping observations.*

Constant-Skew-T (6M)	US	UK	DE	AUS	FR
$\beta_{0,\mu}$	0.08162	0.0767	0.0987	0.0853	0.0954
adj. t-stat	*(5.16)*	*(5.24)*	*(3.40)*	*(5.10)*	*(4.11)*
$\beta_{1,\mu}$	0.0074	0.0108	0.0287	0.0013	0.0203
adj. t-stat	*(0.87)*	*(1.35)*	*(2.32)*	*(0.13)*	*(1.78)*
$\beta_{0,\sigma}$	-2.5807	-2.6426	-2.1176	-2.6262	-2.2286
adj. t-stat	*-(20.96)*	*-(20.86)*	*-(17.84)*	*-(21.59)*	*-(17.31)*
$\beta_{0,\nu}$	-0.2862	-0.2906	-0.2784	-0.2710	-0.2943
adj. t-stat	*-(2.28)*	*-(2.33)*	*-(1.93)*	*-(1.99)*	*-(2.23)*
$\beta_{0,\tau}$	1.3296	1.5104	2.0033	1.5739	1.9211
adj. t-stat	*(3.47)*	*(3.26)*	*(2.90)*	*(3.13)*	*(2.85)*

Table A.5.4: *Illustrates parameter estimates for various international markets using the Constant-Skew-T model for the 6-month holding period. The t-statistics provided in parenthesis are adjusted for overlapping observations.*

Symmetric-T (12M)	US	UK	DE	AUS	FR
$\beta_{0,\mu}$	0.0998	0.0785	0.0889	0.1049	0.0768
adj. t-stat	*(4.88)*	*(3.91)*	*(3.35)*	*(6.08)*	*(2.77)*
$\beta_{1,\mu}$	0.0099	0.0308	0.0793	0.0007	0.0626
adj. t-stat	*(0.55)*	*(1.63)*	*(3.01)*	*(0.04)*	*(2.44)*
$\beta_{0,\sigma}$	-2.1472	-2.0751	-1.7589	-2.2249	-1.6609
adj. t-stat	*-(10.64)*	*-(12.08)*	*-(11.00)*	*-(13.52)*	*-(9.77)*
$\beta_{0,\nu}$	1.1643	2.2267	1.7900	1.5542	3.4540
adj. t-stat	*(2.08)*	*(1.68)*	*(2.30)*	*(2.42)*	*(0.77)*

Table A.5.5: *Illustrates parameter estimates for various international markets using the Symmetric-T model for the 12-month holding period. The t-statistics provided in parenthesis are adjusted for overlapping observations.*

Symmetric-T (6M)	US	UK	DE	AUS	FR
$\beta_{0,\mu}$	0.0500	0.0460	0.0467	0.0569	0.0456
adj. t-stat	*(5.68)*	*(5.59)*	*(3.56)*	*(7.08)*	*(3.71)*
$\beta_{1,\mu}$	0.0087	0.0117	0.0363	-0.0039	0.0285
adj. t-stat	*(1.01)*	*(1.44)*	*(2.82)*	*-(0.38)*	*(2.33)*
$\beta_{0,\sigma}$	-2.5524	-2.6040	-2.0954	-2.6347	-2.1643
adj. t-stat	*-(21.11)*	*-(21.27)*	*-(18.76)*	*-(20.52)*	*-(18.83)*
$\beta_{0,\nu}$	1.2178	1.4089	1.8037	1.3259	1.9342
adj. t-stat	*(3.52)*	*(3.34)*	*(3.28)*	*(3.23)*	*(2.98)*

Table A.5.6: *Illustrates parameter estimates for various international markets using the Symmetric-T model for the 6-month holding period. The t-statistics provided in parenthesis are adjusted for overlapping observations.*

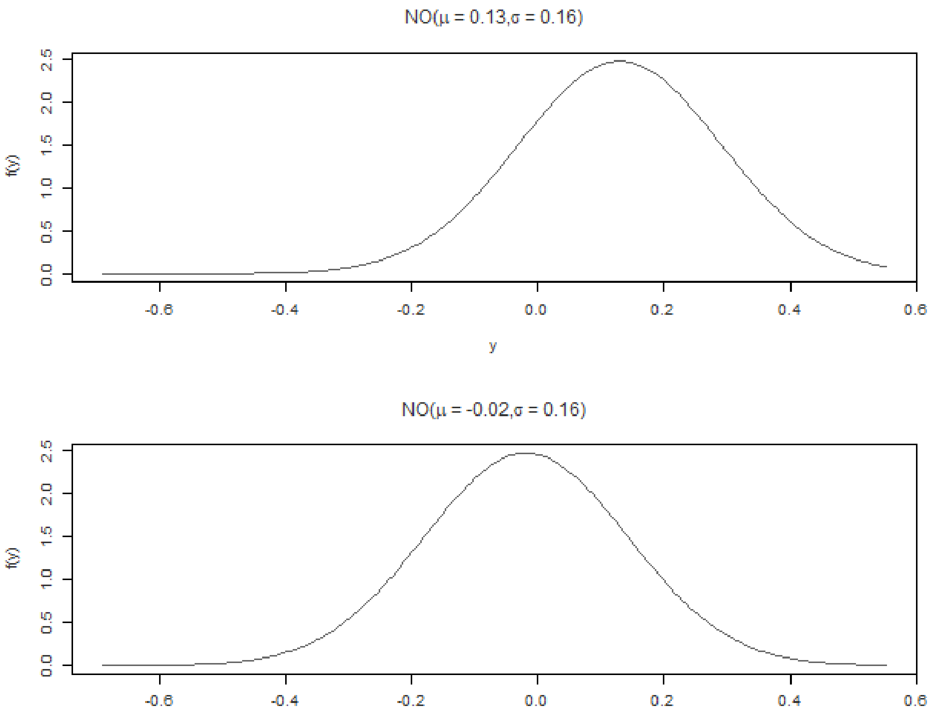

Figure A.5.7: *12M - The figure depicts the Normal model's implied 12-month excess return conditional on the slope factor. The top graph is implied excess returns conditional on high (2 standard deviations above the mean) slope readings, whereas the bottom graph are implied excess returns conditional on low (2 standard deviations below the mean) slope readings.*

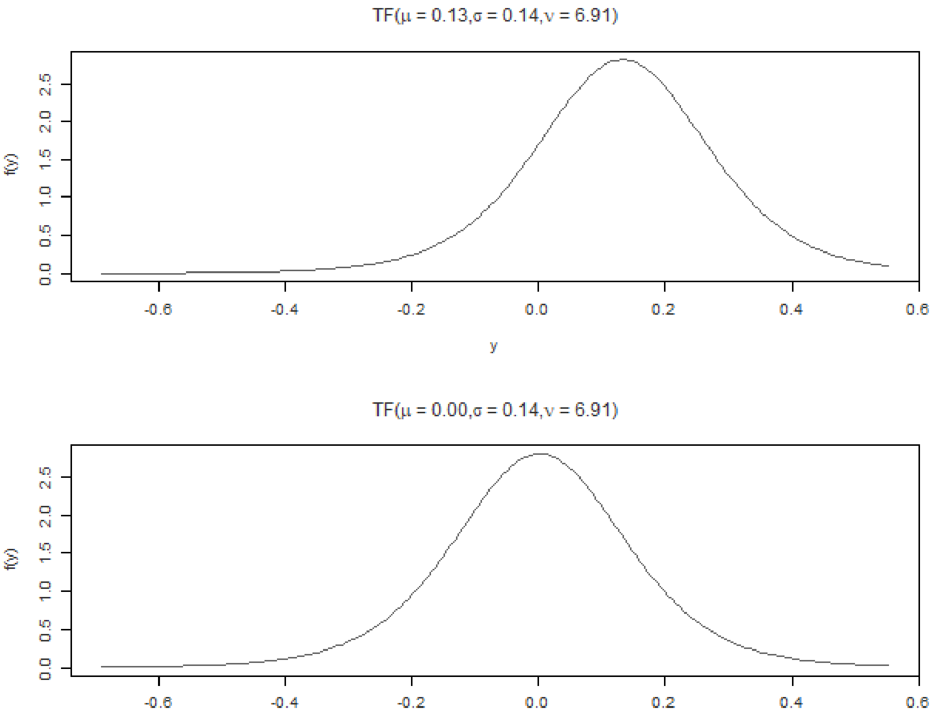

Figure A.5.8: *12M - The figure depicts the Symmetric-T model's implied 12-month excess return conditional on the slope factor. The top graph is implied excess returns conditional on high (2 standard deviations above the mean) slope readings, whereas the bottom graph are implied excess returns conditional on low (2 standard deviations below the mean) slope readings.*

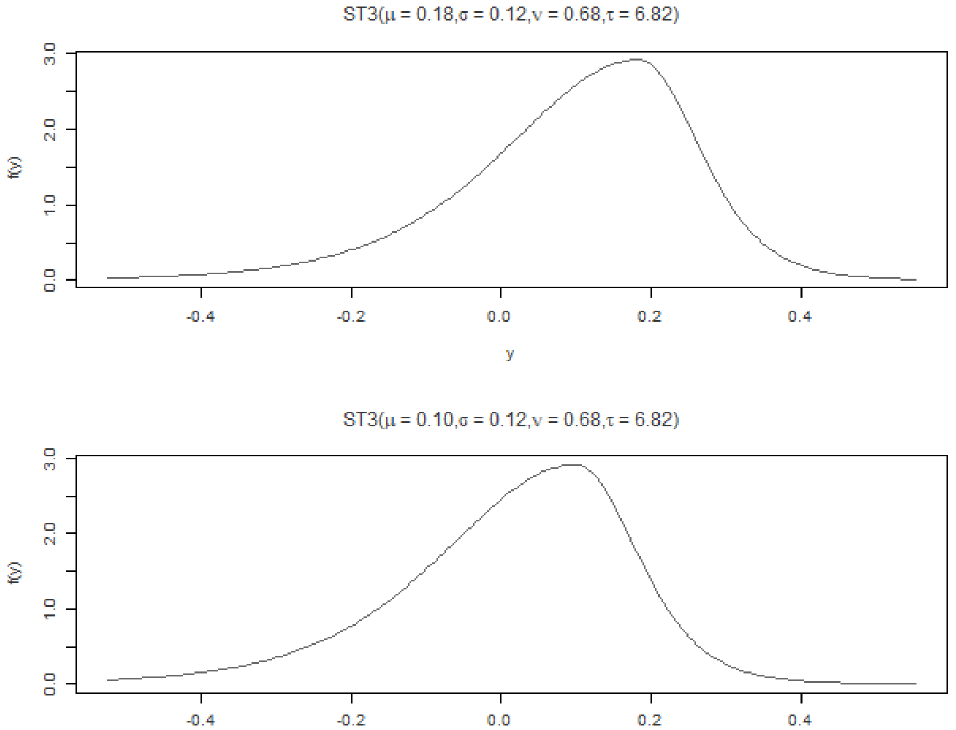

Figure A.5.9: *12M - The figure depicts the Constant-Skew-T model's implied 12-month excess return conditional on the slope factor. The top graph is implied excess returns conditional on high (2 standard deviations above the mean) slope readings, whereas the bottom graph are implied excess returns conditional on low (2 standard deviations below the mean) slope readings.*

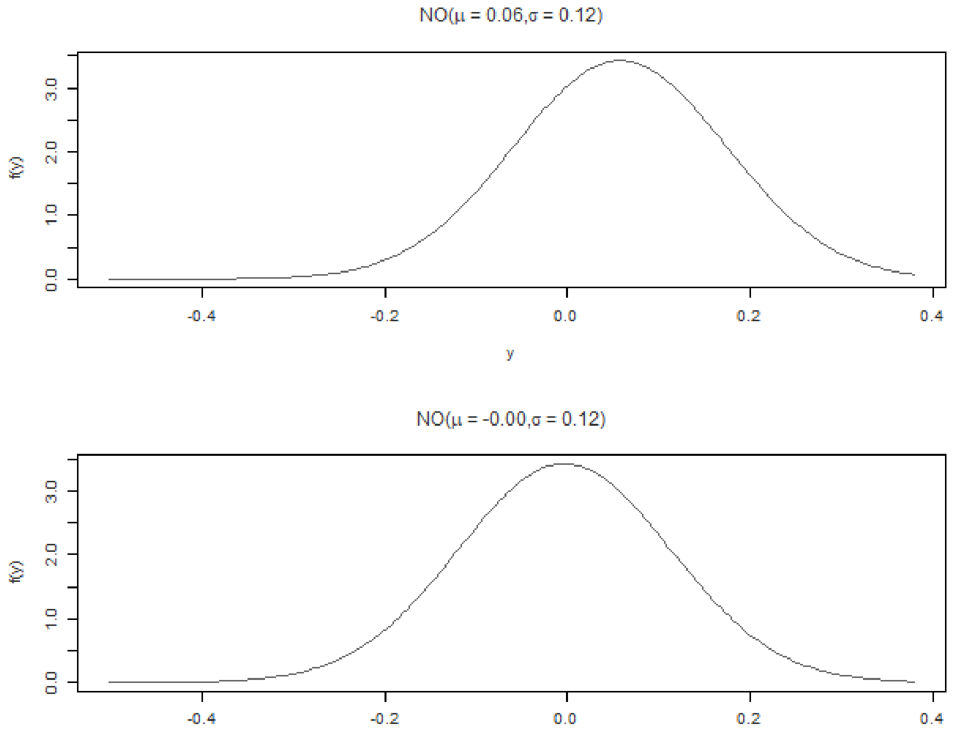

Figure A.5.10: *6M - The figure depicts the Normal model's implied 6-month excess return conditional on the slope factor. The top graph is implied excess returns conditional on high (2 standard deviations above the mean) slope readings, whereas the bottom graph are implied excess returns conditional on low (2 standard deviations below the mean) slope readings.*

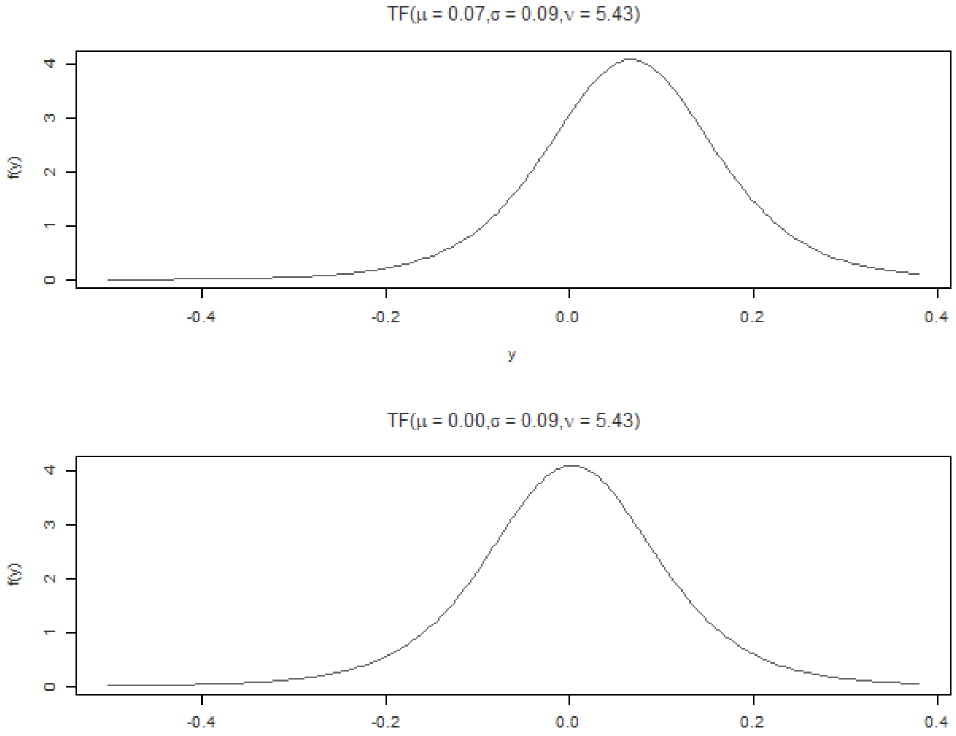

Figure A.5.11: *6M - The figure depicts the Symmetric-T model's implied 6-month excess return conditional on the slope factor. The top graph is implied excess returns conditional on high (2 standard deviations above the mean) slope readings, whereas the bottom graph are implied excess returns conditional on low (2 standard deviations below the mean) slope readings.*

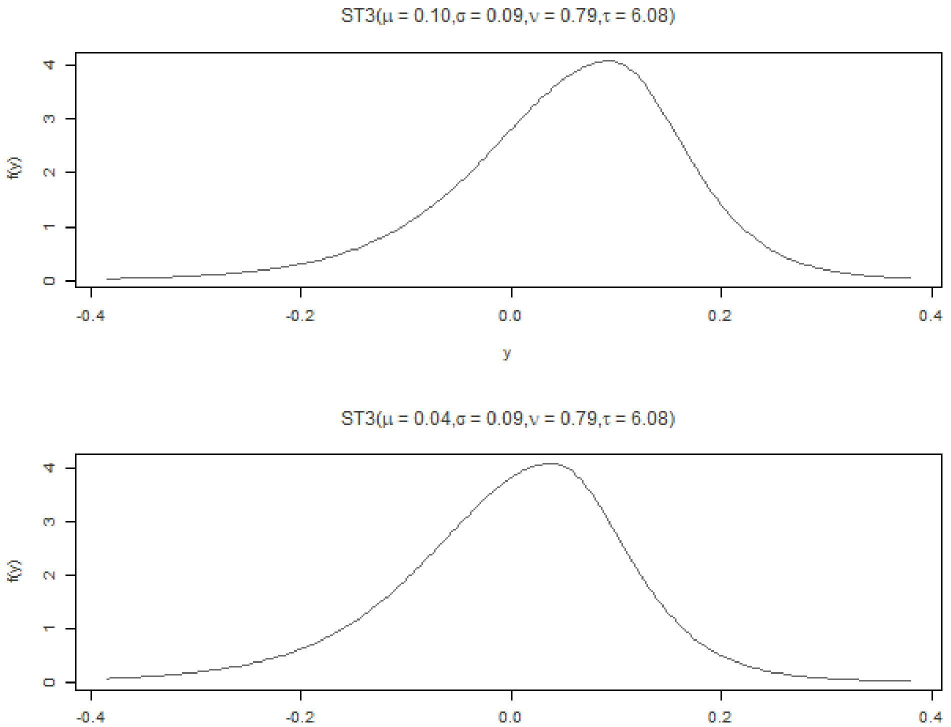

Figure A.5.12: *6M - The figure depicts the Constant-Skew-T model's implied 6-month excess return conditional on the slope factor. The top graph is implied excess returns conditional on high (2 standard deviations above the mean) slope readings, whereas the bottom graph are implied excess returns conditional on low (2 standard deviations below the mean) slope readings.*