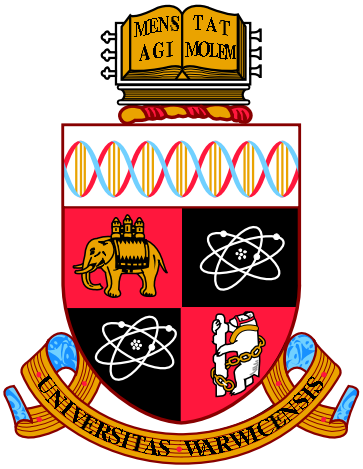

Essays on Asset Pricing and Macro-Finance

by

Jingwen Shi

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Warwick Business School

September 2019

Contents

List of Tables

List of Figures

Acknowledgments

I am deeply grateful to my supervisor Prof. Shaun Vahey, an outstanding scholar with thorough knowledge and expertise, a caring person with genuine sympathy and compassion, and an inspirational mentor with great wisdom and insight. The most important lesson he has taught me is how to remain a beacon of positivity in the face of sufferings and challenges, which I will always carry with me on the way forward.

My sincere gratitude goes to Prof. Anthony Garratt, one of the nicest people imaginable, for his constant support and valuable suggestions along the way. It's the warm kindness I receive from him that will inspire me to do good things for others.

My deep appreciation goes to Prof. Francesco Ravazzolo for generously offering his support, guidance and inspiration to help me complete my PhD. I would like to express my special thanks to Dr Domenico Giannone for his constructive comments and insightful advice during the course of this thesis. I would also like to thank Prof. Xing Jin for his supervision and guidance, Prof. Philippe Mueller for his helpful comments in my reviews, Dr Elizabeth Whalley for helping me get back on track, Prof. Michael Moore for admitting me into the programme, and many staff in Warwick Business School with whom I had meaningful conversations with.

Many thanks go to my colleagues and friends, from whom I learn a great deal and with whom I share lots of joy and pain.

My endless gratitude goes to my parents. Their being there for me is what sustains me through all arduous moments, gives me the strength to fight another day, and motivates me to keep moving forward.

Declarations

I declare that this thesis has not been submitted, in whole or in part, for any previous degree. The work contained herein is my own except where explicitly stated otherwise by reference or acknowledgement. I further declare that Chapter 4 is my job market paper.

Abstract

This thesis evaluates and proposes forecasting methods for two primary categories of investments, stocks and bonds, aiming to advance our understanding of the underlying risk–return attributes and to explore the scope for enhanced predictive content in practice. More specifically, the thesis identifies and examines the cross-sectional and time-series determinants of stock returns and government bond yields from a predictive perspective.

Chapter 2 examines the information content of variance swaps in predicting the cross section of stock returns. A non-parametric approach is employed to extract principal components as underlying predictive variables, which mitigates issues concerning model misspecification and estimation uncertainty with parametric approaches. The variance swap components exhibit substantial predictive power for the cross-sectional stock returns. The last few principal components, despite carrying little weight in terms of explaining the variation in the variance swap surface, contribute considerably towards predicting stock returns. The information value of variance swaps is attributable to their strong associations with financial and macroeconomic state variables that track investment opportunities, and with various moment measures that characterise the variations in aggregate stock market returns. Providing additional information beyond that afforded by benchmark factors, the variance swap components, combined with Fama–French–Carhart factors, deliver superior forecasts.

Chapter 3 evaluates the predictability in government bond markets implied by yield forecasting approaches that exploit macro-finance interactions in a sample extended to the post-global financial crisis period, which spans 1970:M1–2016:M12. This chapter places particular focus on the macro-yields model featuring unspanning restrictions proposed by Coroneo, Giannone, and Modugno (2016), which has demonstrated satisfactory performance in forecasting the yield curve and excess bond returns within the 1970:M1–2008:M12 timespan. While the model's in-sample fit alters little, the out-of-sample predictability deteriorates substantially as a result of severe underprediction of all three Nelson-Siegel yield factors during the 2007–09 global financial crisis (GFC), and the overprediction of yields in the subsequent recovery. The chapter further identifies the leading causes of the deterioration and investigates the stability of yield dynamics and macro-financial linkages, by constructing conditional forecasts and counterfactual scenarios. Even after controlling

for the anomalous behaviour of the front-end of the curve, the predictability of longer government bonds still suffers from substantial declines. Meanwhile, the instability of yield dynamics and macro-finance interdependence is revealed in the context of the linear-Gaussian dynamic factor representation. These results collectively suggest that the severe model degradation cannot be solely attributed to the constraints on the dynamics of the short-term rates. It also emphasises the importance of addressing the growing complexity in yield dynamics and macro-finance interactions after the global financial crisis, which cannot be adequately accommodated by the existing macro-finance term structure framework.

Chapter 4 explores the potential of the copula framework to handle non-Gaussianity in the context of macro-finance term structure modelling and forecasting. The non-Gaussian macro-yields model proposed here accounts for the asymmetry and tailedness in yield distributions through non-parametric marginal densities, as well as explicitly addressing the cross-sectional and serial dependence via the state-space inversion copula, and in doing so retains a latent dynamic factor structure amenable to efficient implementation. Regardless of maturities and forecast horizons, exploiting the informational content of macroeconomic data in a non-Gaussian setting improves both in-sample and out-of-sample forecasting performance relative to the Gaussian macro-yields model over the 1970:M1–2016:M12 period. Furthermore, the non-Gaussian macro-yields model demonstrates overwhelming superiority in predicting excess bond returns over the expectations hypothesis and the prominent macro-financial predictors. It also compares favourably with the random walk in forecasting the yield curve over medium- to long-term horizons. Lastly, this copula-based approach affords a technically convenient and extensible means of accommodating high-dimensional macroeconomic datasets and the growing complexity of post-crisis yield movements, which facilitates further investigation into their practical implications for government bond modelling and forecasting.

Abbreviations

CP	Cochrane–Piazzesi
DFM	Dynamic Factor Model
DNS	Dynamic Nelson–Siegel
FF3	Fama–French Three-Factor
FFC	Fama–French–Carhart
GFC	Global Financial Crisis
HML	High-Minus-Low (Value)
IC	Information Criterion
LN	Ludvigson–Ng
MSFE	Mean Square Forecast Error
PCA	Principal Component Analysis
SMB	Small-Minus-Big (Size)
UMD	Up-Minus-Down (Momentum)
VAR	Vector Autoregression
VIX	The Volatility Index
VS	Variance Swap

Chapter 1

Introduction

The risk–return profiles of asset classes contain valuable information regarding the aggregate perceptions of market participants. Active investors exploit these to construct profitable positions based on their anticipation of market changes, and policymakers examine these to seek appropriate ways to promote overall stability and well-being. This thesis investigates the determinants underlying the time-series and cross-sectional variations in stock returns and in government bond yields, particularly from a predictive perspective. Investigating the nature of these factors is instrumental in developing reliable valuation and prediction models that adequately capture the latent structure, facilitating investment decisions and inferences. The empirical performance of these models sheds further light on the distributional characteristics and evolutionary patterns of the latent states.

Central to the factor modelling approach is the notion that the dynamics of a relatively large collection of observables are typically governed by a smaller set of latent variables. Analogous to the financial argument that underlying the pricing of numerous financial instruments are a small number of systematic risk factors, econometric studies indicate that driving the evolution of various economic variables is a common set of fundamental forces. This study investigates the implications of these research findings across stocks with different firm characteristics in Chapter 2, and across government bonds of various maturities in Chapter 3 and Chapter 4.

The factor modelling approach to asset pricing links pricing kernels that produce contingent claim valuations to observable or latent risk factors, interpreting excess returns as compensation for the corresponding risks. In a multi-period context, investors also take into account the risks associated with future shifts in consumption and investment opportunities when determining the cross-sectional risk–return trade-off. The term structure of variance swaps represents an important source of information regarding the risk-neutral expectations of future market

volatility over different horizons. Inspired by this, Chapter 2 examines and interprets the information content of variance swaps in the context of forecasting the cross section of stock returns. Providing additional information beyond that afforded by benchmark factors, the combination of variance swap components and Fama–French–Carhart factors achieves enhanced forecasts for cross-sectional stock returns. The information content of these components is attributed to their associations with state variables that track changes in the investment opportunity set and moment measures that characterise overall market conditions.

Constituting another major asset category, bonds represent relatively conservative investments with steady income streams and less volatile returns, making them an available counterbalance to investors' equity holdings. During times of turbulence and uncertainty when market sentiment leans towards the negative, investors are inclined to seek the relative safety of bonds, and thus capital flows into fixed-income markets. Such perceptions are a key ingredient in policymaking that determines whether to boost aggregate demand via investment activity. Also incorporated into monetary policy decision-making are a number of indicators that assess the current economic situation, especially those concerned with real economic activity and inflation, linking short-term rates directly to these macroeconomic aspects. Accordingly, with regard to the inherent risks, bonds are generally less vulnerable to price swings than equities but have greater susceptibility to uncertainty in terms of monetary policy and the dynamics of the broader economy. The yield curve, with the shorter-end directly influenced by the monetary authority and the longer-end reflecting market participants' expectations, represents an insightful tool to elucidate interactions between macroeconomic fundamentals and financial markets.

The factor structure has also demonstrated its capability to provide a highly accurate characterisation of government bond yields across the maturity spectrum, besides its applications in the equity market. The dynamic factor approach employs vector autoregressive (VAR) processes to model the latent state variables, whereby the dynamics and interactions translate into cross-serial and cross-sectional dependencies among the observable variables. Within this framework, the following chapter of this thesis examines and quantifies the changes in predictive power of macroeconomic variables for government bond yields in the aftermath of the global financial crisis.

Chapter 3 evaluates the predictability in government bond markets over a prolonged period that extends to the post-global financial crisis era in the context of the macro-yields model with unspanning properties (Coroneo et al., 2016). Having produced satisfactory outcomes over the 1970:M1–2008:M12 period, the macro-finance modelling suffers substantial deterioration in terms of both in-sample and

out-of-sample forecast performance in the extended 1970:M1–2016:M12 period, no longer presenting conclusive evidence in favour of the macro-finance framework. Further statistical analysis attributes the general decline in predictability to shifts in the model-implied dependence structure and furthermore to the practical limitations of the linear-Gaussian framework, particularly during turbulent periods. This chapter also considers the use of real-time data and assesses the impact of revisions. Caruso and Coroneo (2019) subsequently develop this direction by employing a mixed-frequency model that incorporates quarterly interest rate surveys, and arrive at similar conclusions.

In the context of the dynamic factor representation, Chapter 3 investigates the stability of yield dynamics and macro-finance interactions during and after the crisis. In particular, this chapter examines whether the degradation in performance across the entire yield curve spectrum stems from possible instabilities over the front-end of the curve. To control for the anomalous behaviour of short-term rates, this chapter exploits the conditional forecast technique to construct and evaluate counterfactual scenarios. Certain maturity spectra of the curve are considered unobservable post-2008:M10 and imputed by projections of the realised paths of other variables according to pre-crisis patterns and structures (Bańbura et al., 2015; Giannone et al., 2019; McCracken and McGillicuddy, 2019). While the in-sample conditional forecasts reveal notable discrepancies between counterfactual paths and actual values, the conditional forecasting approach offers no satisfactory solution to the out-of-sample forecasting problem over this extended sample.

Collectively, these results imply that the existing macro-finance modelling framework is incapable of characterising the post-crisis yield dynamics and macro-yields interdependence. This highlights the importance of accounting for the growing complexity of yield marginals and macro-financial linkages in the low-interest-rate environment. To provide a more plausible and robust means of applying the factor structure in the post-crisis period, Chapter 4 seeks to generalise the macro-yields model examined in Chapter 3 by relaxing its commonly held assumption of Gaussianity. Drawing upon the insights of Smith and Vahey (2016), which consider the asymmetry in macroeconomic modelling, this chapter proposes a non-Gaussian macro-yields model that plausibly accommodates the univariate features observed in yields (e.g. asymmetry, heavy-tailedness, and multi-modality) via fitting asymmetrical non-parametric marginal distributions while explicitly addressing the cross-sectional and serial dependence by means of the state-space inversion copula. In doing so, the proposed model can be regarded as a generalisation of the original macro-yields model, encompassing the latter as a special case that specifies normal marginal distributions for all variables.

Over the 1970:M1-2016:M12 period, this non-Gaussian macro-yields model delivers superior forecasts compared to its Gaussian rival in all instances regardless of maturities and forecast horizons, as indicated by an average reduction of 0.39 in the relative mean squared forecast errors (rMSFEs) for the yield curve and 0.59 for excess bond returns. Even when compared with the no-predictability benchmarks, the proposed model demonstrates competitive advantage, in particular for medium- to long-term forecast horizons. Distinguished from the Gaussian setting in which the marginal and forecast densities are constrained to be symmetric and unimodal, the copula approach plausibly captures the multimodal and asymmetric aspect of the distributions of bond yields in the post-crisis period.

This study has gone some way to furthering our understanding of the predictive mechanisms underlying the macro-finance modelling, and presenting possible solutions to the challenges posed by the recent recession. By additionally imposing a dynamic factor structure for dimensionality reduction, this study represents a productive attempt to explore the usefulness of the state-space inversion copula formulation in dealing with high-dimensional datasets in the macro-finance context. In doing so, this study provides an empirical avenue for exploring the potential of the copulas in the vast field of macro-finance research in a data-rich environment. The implicit copula approach explored here can conveniently retrieve inherent structures from any established multivariate representations without being limited to existing canonical distributions, contributing towards a more comprehensive characterisation of the interactions between the macroeconomy and asset markets.

Chapter 2

The Information Content of Variance Swaps in Predicting the Cross Section of Stock Returns[1]

2.1 Introduction

Stock return predictability has long been an issue of significance for practitioners and academics, which has implications for investment decision making and general equilibrium modelling. Previous research has identified a wide range of factors that explain and predict stock returns, such as financial valuation ratios, macroeconomic variables, and other market indicators. Researchers are becoming increasingly aware of the critical roles of variance risk in determining asset returns (Bakshi and Kapadia, 2003; Ang et al., 2006; Adrian and Rosenberg, 2008; Cremers et al., 2015) and of variance risk premium in predicting aggregate market returns (Bollerslev et al., 2009; Drechsler and Yaron, 2011). Accordingly, appropriate characterisation of the underlying variance process and accurate estimation of the associated risk premium are attracting growing interest, along with the development of many variance-related derivatives over the past few decades. The variance risk premium, reflected in the difference between the expected realised variance under physical and risk-neutral measures, is widely-acknowledged to contain rich information regarding the overall economic uncertainty and aggregate risk aversion. The existing literature seeks to extract implied distributions and time-varying risk premia from a range of financial instruments, such as the aggregate market index; index options; the Volatility Index (VIX) and derivatives written on it, such as VIX futures and VIX options; and the most straightforward and actively-traded variance derivative, the variance swaps.

To retrieve the extensive information embedded in variance swaps, previous studies propose parametric models to accommodate principal stylised features of the

[1]This chapter was conducted under the supervision of Prof. Xing Jin during his tenure at Warwick Business School.

term structure of variance swaps. These approaches require a well-identified set of underlying states and well-specified assumptions within a random field framework in order to maintain tractability. Focusing on the information content of variance swaps without a cumbersome number of distributional assumptions on the underlying process and relevant parameters, this chapter employs a non-parametric approach by utilising the principal components to extract information pertinent to the time-series and cross-sectional determinants of stock returns. Principal component analysis is frequently employed to provide a general description of the cross section of the variance swap itself, through which the number of underlying states is determined. The implications of these components in the cross section of stock returns, particularly from a predictive perspective, have not been fully explored in literature to date. While closely related to the work by Hong and Jin (2018), which forwards the time-series predictability of variance swaps, this chapter extends the scope into cross-sectional analysis and traces the sources of predictability in this context. Thus, this chapter adds to the growing body of literature on variance swaps by evaluating the information content of the variance swap components, constructed directly from market quotes utilising a non-parametric approach, in the cross section of stock returns.

Presenting several advantages over parametric methods, the non-parametric approach circumvents some of the issues associated with the modelling and estimation procedures. Without predetermining the number of underlying factors, the non-parametric approach addresses concerns about model misspecification to some extent. Current research typically focuses on examining the first three principal components, which together explain over 90% of the variation in variance swap curves. However, the remaining few components, despite carrying negligible roles in representing the variation in variance swaps, hold considerable potential for forecasting cross-sectional stock returns. Allowing more components to be exploited, the non-parametric approach enables incorporation of more comprehensive information that cannot be adequately captured by model-based approaches that summarise the entire surface into two to three latent states. The rapid implementation of principal component analysis avoids limitations of the extremely time-consuming and cumbersome procedures of the combined state-parameter estimation, as well as the additional uncertainty induced by the estimation algorithms. Moreover, by constructing factors directly from market quotes, this chapter implicitly incorporates information concerning not only the underlying states, but also the time-varying risk premia.

Characterising the cross-sectional variation of equity returns has been an important focus of asset pricing research for decades. Besides the benchmarks based

on the characteristics of stocks, the literature also considers statistical measures that gauge variation in market conditions and serve as indicators of the market-wide risks. Analogous to implied ex-ante measures derived from option prices, variance swaps, constructed from expectations about future market variation, also represent a means of retrieving the "forward-looking" dynamics and distributions. This work is empirically inspired by the notion of "ex-ante assessment" of aggregate volatility and jump risk and theoretically grounded on parametric models that relate the term structure of variance swaps to these latent states (Amengual, 2009; Egloff et al., 2010; Filipović et al., 2016; Li and Zinna, 2018; Ait-Sahalia et al., 2018).

Bridging the literature on the cross-sectional pricing and time-series predictive implications of variance-related factors, this research evaluates the predictive content of the variance swaps in the cross section of stock returns. The variance swaps contain substantial predictive information regarding the cross-sectional stock returns, beyond those afforded by the benchmark Fama–French–Carhart factors that capture the size, value and momentum effects. Therefore, combining the variance swap components with the Fama–French–Carhart factors enables enhanced forecasting of excess stock returns at all horizons.

Probing into sources of the predictability, this chapter provides further insight into the associations between the variance swap components and the prevalent state variables and various moment-based measures. Variance swap components exhibit significant and substantial explanatory and predictive power for the state variables that correspond to business-cycle fluctuations and the statistical measures that depict conditional distribution of aggregate market returns. Providing additional information while not fully spanning the benchmark factors, variance swap components, combined with Fama–French–Carhart factors, achieve enhanced performance in terms of predicting the cross section of stock returns.

The rest of this chapter is structured as follows. Section 2.2 reviews the related research, with Section 2.3 describing the empirical framework. Section 2.4 presents the results of the predictability in the cross-sectional stock returns, before Section 2.5 proceeds by interpreting the information content in variance swaps and Section 2.6 concludes the chapter.

2.2 Related Literature

It has been widely acknowledged that the prices of various derivatives contain rich information regarding the underlying volatility process. Research has been conducted to recover the implied distributions from various market observations, including the aggregate market index (Eraker et al., 2003; Jacod et al., 2010; Fulop et al., 2014);

index options (Eraker, 2004; Broadie et al., 2007); VIX and related derivatives, such as VIX futures and VIX options (Zhang and Zhu, 2006; Dotsis et al., 2007; Lin, 2007; Zhu and Zhang, 2007; Wang and Daigler, 2011); and variance swaps. Besides the close association with the underlying volatility process, variance swaps present another advantage by admitting a straightforward analytical structure in the latent variables, which could be affine, quadratic or polynomial depending on the model specifications, thus enabling the risk premium to be directly recovered from the variance swap rates. Despite providing investors with the most direct exposure to variance risk and serving as beneficial tools in the context of risk management and asset allocation, variance swap contracts are not as readily accessible as other derivatives such as options. It is only during the past decade that researchers have begun to exploit variance swap quotes to retrieve information about the market variation (Egloff et al., 2010). The majority of studies focus on proposing parametric stochastic models that extend the traditional stochastic volatility framework, prominent among which are the multivariate pure diffusion models examined in Egloff, Leippold, and Wu (2010) and Filipović, Gourier, and Mancini (2016), and the jump-related models studied in Amengual (2009), Ait-Sahalia, Karaman, and Mancini (2018), and Li and Zinna (2018).

An additional strand of the literature closely related to this chapter is the investigation into the relationships between stock returns and volatility and jump risks from two perspectives: the pricing implications in cross-sectional stock returns and the predictive implications in aggregate market returns. In terms of cross-sectional analysis, it has been demonstrated that market volatility is significantly priced in the cross section of stock returns (Ang et al., 2006). Further disentangling market volatility into long-run and short-run components generates greater explanatory power along the cross-sectional dimension (Adrian and Rosenberg, 2008). From a predictive point of view, distinguishable from conventional predictors, the variance risk premium provides a new source of predictability and displays substantial predictive power for short-horizon returns (Bollerslev et al., 2009, 2011; Drechsler and Yaron, 2011). Similar analyses are carried out in foreign stock markets (Londono, 2015), the credit market (Buraschi et al., 2013), and the international market from a global perspective (Bollerslev et al., 2014).

Separating jumps from path-wise variations, Cremers, Halling, and Weinbaum (2015) examine the pricing of aggregate jump risk in the cross section of expected returns. Also, jump tail risk has been demonstrated to enhance the predictability afforded by the variance risk premium, while it plays different roles from the diffusive risk in terms of return horizons and portfolio construction (Li and Zinna, 2018). The current chapter builds upon this strand of literature by incorpo-

rating both variance and jump aspects and the corresponding risk premium information through the variance swap components when examining the cross-sectional predictability. Jump-diffusion parametric models provide solid theoretical foundations that associate variance swaps with the volatility and jump risks (Amengual, 2009; Ait-Sahalia et al., 2018; Hong and Jin, 2018; Li and Zinna, 2018).

This chapter further examines the information content of variance swaps by directly relating them to two important categories of observable equity factors. The literature has identified a broad range of variables as predictors of the equity premium (Petkova, 2006; Welch and Goyal, 2007; Pettenuzzo et al., 2014), such as dividend yields (Fama and French, 1988), the Treasury yields and term spread that capture the variations in the bond market (Patelis, 1997; Hahn and Lee, 2006), and default spread that relates to the conditional distribution of asset returns (Campbell and Thompson, 2007). From a different viewpoint, the moment measures of return distributions are also considered; for example, Chabi-Yo (2012) introduces stochastic volatility, stochastic skewness, and stochastic kurtosis into the pricing kernel. In this context, option prices are extensively explored to provide estimates of the implied skewness (Cremers and Weinbaum, 2010; Xing et al., 2010; Conrad et al., 2013; Bali and Murray, 2013). This chapter examines the associations between the variance swap components and prevailing equity factors as well as with various moment measures in order to provide certain insights into this vein of interest.

2.3 Predictability in the Cross Section of Stock Returns

This chapter places particular focus on examining the information content of variance swaps on the five-by-five double-sorted quintile portfolios formed on size and book-to-market ratio, i.e. 25 double-sorted portfolios (Fama and French, 1992, 1993, 1996). This section briefly discusses the conceptual framework and empirical approaches used to conduct predictive analytics in the cross-sectional dimension.

2.3.1 Theoretical Motivations

This subsection reviews the parametric models proposed in Egloff, Leippold, and Wu (2010), Ait-Sahalia, Karaman, and Mancini (2018) and Hong and Jin (2018) to provide insights into the information content of variance swaps from a model-based perspective. Applying no-arbitrage arguments, the variance swap rate is determined in terms of the expected values of aggregate variance under the risk-neutral measure. Given a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, P)$, a general semi-martingale

setting is typically adopted for the index process

$$dS_t/S_{t_-} = (r-d)dt + \sqrt{(1-\rho^2)v_t}dW_{1t}^Q + \rho\sqrt{v_t}dW_{2t}^Q + [\exp(J_t^{s,Q}) - 1]dN_t - g^Q\lambda_t dt, \tag{2.1}$$

with r referring to the risk-free rate, d to the dividend rate, and v_t to the instantaneous variance, which follows the Cox-Ingersoll-Ross process (Amengual, 2009; Egloff et al., 2010) with central tendency m_t and jumps $J_t^{v,Q}$

$$dv_t = \kappa_v^Q(m_t - v_t)dt + \sigma_v\sqrt{v_t}dW_{2t}^Q + J_t^{v,Q}dN_t, \tag{2.2}$$

$$dm_t = \kappa_m^Q(\theta_m^Q - m_t)dt + \sigma_m\sqrt{m_t}dW_{3t}^Q. \tag{2.3}$$

The diffusive movements are driven by independent Brownian motions W_{it} and the jump component is characterised by the counting process N_t with size parameters J_t^i. The leverage effect, ρ in Equation (2.1), is included to account for the negative interactions between returns and changes of volatility (Carr and Wu, 2017). Through this specification, the prices of variance swaps can be easily obtained from the quadratic variation of the log-price (Hong and Jin, 2018)

$$\begin{aligned} VS_{t,t+\tau} =& \bar{v}_{t,t+\tau}^Q + E_t^Q[(J^s)^2]\bar{\lambda}_{t,t+\tau}^Q \\ =& \frac{1}{\tau}\int_t^{t+\tau} E_t^Q[v_s]ds + \frac{1}{\tau}E_t^Q[(J^s)^2]\int_t^{t+\tau} E_t^Q[\lambda_s]ds, \end{aligned} \tag{2.4}$$

with λ denoting the intensity of the counting process.

The affine framework (Duffie et al., 2000) yields straightforward expressions for variance swap rates in terms of the underlying states, enabling the latent factors to be conveniently recovered from variance swap quotes. Setting jump parameters to zero produces the pure-diffusion model in Egloff, Leippold, and Wu (2010), while assuming a linear relationship between λ and v gives the model presented in Ait-Sahalia, Karaman, and Mancini (2018), both of which conform to

$$VS_{t,t+\tau} = a(\tau;\Theta) + b(\tau;\Theta)'[v_t, m_t]', \tag{2.5}$$

where $a(\tau;\Theta)$, $b(\tau;\Theta)$ are functions of parameters that depict the risk-neutral dynamics. Hong and Jin (2018) specify an independent process for λ_t, which results in

$$VS_{t,t+\tau} = a(\tau;\Theta) + b(\tau;\Theta)'[v_t, m_t, \lambda_t]', \tag{2.6}$$

or more specifically,

$$VS_{t,t+\tau} = \frac{1}{\tau}\int_t^{t+\tau} E_t^Q[v_s]ds + \frac{1}{\tau}E_t^Q[(J^s)^2]\int_t^{t+\tau} E_t^Q[\lambda_s]ds \tag{2.7}$$

$$= \phi_\theta(\tau)\theta_m^Q + \phi_{\lambda_\infty}(\tau)\lambda_\infty + \phi_v(\tau)v_t + \phi_m(\tau)m_t + \phi_\lambda(\tau)\lambda_t, \tag{2.8}$$

where ϕ are deterministic functions of the risk-neutral parameters, which are related to their physical counterparts via the market prices of risks.

As demonstrated by Equations (2.5–2.8), the term structure of variance swaps is explicitly associated with the underlying states, including the instantaneous variance v_t, the central tendency m_t and the jump intensity λ_t, which represent the short- and long-run components of market volatility and the jump factor, respectively. Moreover, the term structure of variance swaps contains information regarding the risk premia that compensate investors for bearing corresponding sources of uncertainty. Thus, compared with a parametric approach that extracts the underlying states v_t, m_t, λ_t as latent factors, variance swap components constructed directly from market quotes incorporate additional information on the risk premia, thus revealing insights into the overall economic uncertainty and aggregate risk aversion.

2.3.2 Empirical Framework

To reconcile time-series and cross-section analysis, this chapter applies the traditional Fama–MacBeth approach (Fama and MacBeth, 1973) that implements a two-pass procedure. As outlined by Bali, Engle, and Murray (2016), the cross-sectional dispersion is accommodated by

$$r_{i,t+h} = \delta_{0,t} + \delta_{M,t}\beta_{i,M} + \sum \delta_{K,t}\beta_{i,K}, \tag{2.9}$$

where $r_{i,t+h}$ refers to the h-period-ahead excess return of each asset i, while $\beta_{i,M}$ and $\beta_{i,K}$ gauge asset sensitivities to the corresponding factors

$$r_{i,t} = \alpha_i + \beta_{i,M}F_{M,t} + \sum \beta_{i,K}F_{K,t} + \epsilon_{i,t}, \tag{2.10}$$

with $\epsilon_{i,t}$ denoting the regression residuals. Time series regressions in Equation (2.10) are first performed to measure factor exposures β, based upon which the relationship between cross-sectional returns and factor loadings is examined by Equation (2.9).

This chapter focuses on exploring the predictive content of variance swaps in the 5x5 portfolios formed by assets sorted on the basis of different size and value attributes. The selection of characteristic-sorted portfolios as base assets facilitates direct comparison of candidate predictors with the Fama–French benchmark and

other established predictors. The individual characteristics are either accounted for by the size (small-minus-big [SMB]) and value (high-minus-low [HML]) factors, or averaged out in these portfolios. Using a fixed forecasting scheme, the entire period is split into two sub-samples, utilised for estimating the beta representation in the time-series dimension (Equation (2.10)) and for evaluating the predictability in the cross-sectional dimension (Equation (2.9)), respectively. With a rolling window scheme, time-varying betas are sequentially computed using daily observations of the 12-month period prior to and including period t, based on which the cross-sectional studies are conducted.

While the typical cross-sectional studies are mainly concerned with the contemporaneous effects, this research assesses the cross-sectional associations from a predictive perspective and thus uses future excess returns $r_{t,t+h}$ in the periodic cross-sectional regressions, where h (=1M, 3M, 6M, 12M, 24M) indicates h periods ahead. In doing so, this chapter utilises 1-, 3-, 6-, 12-, and 24-month future excess returns as dependent variables in order to examine the cross-sectional predictability over various horizons. To enable direct comparison with benchmark models, the above procedures are implemented with a range of specifications that exploit the market factor, the size and value factors, the moment factor, other predictor variables and the corresponding extended set of variables augmented with variance swap components.

2.4 Empirical Analysis

The first subsection provides a general description of the variance swap term structure and selects the optimal number of components according to the information criterion, based on which the predictability is assessed in the following subsection.

2.4.1 Summary Statistics

This chapter utilises daily variance swap rates on the S&P 500 with the maturity of 1–6-, 9-, 12-, 18- and 24-month periods spanning from November 2008 to September 2017. The 25 size and book-to-market sorted portfolio returns are available from the Kenneth R. French Data Library; while the market factor (Mkt), size factor (SMB), value factor (HML), and momentum factor (UMD) are collected from the Centre for Research in Security Prices (CRSP). The index series spans the highly volatile 2007–09 financial crisis and the European debt crisis. The market turmoil creates considerable turbulence that cannot be accommodated by the diffusive movement, as suggested by the dramatic fluctuations in the index process (upper panel in Figure 2.1) and the sudden bursts of volatility (middle panel in Figure 2.1), indicating the

presence of jumps. According to the negative skewness (-0.475) and high kurtosis (10.647) displayed in Table 2.1, index returns are skewed left, featuring a longer tail on the left side, and leptokurtic, implying fat-tailed distributions.

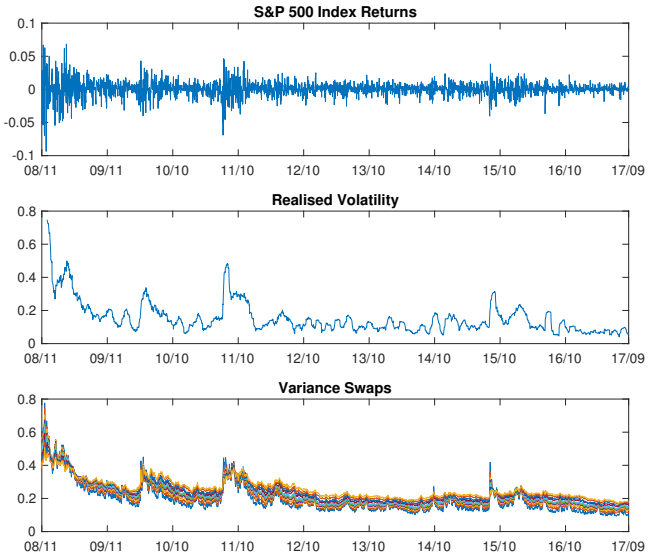

Figure 2.1: Index returns, realised volatility, and variance swap rates. The figure shows the S&P 500 index returns (top panel), realised volatility (middle panel), and variance swap rates (bottom panel) from November 2008 to September 2017. Realised variance is calculated by taking the monthly sum of squared daily returns and is expressed in annualised volatility terms.

Table 2.1: Descriptive statistics for index returns and variance swap rates

	SPX	VS_1M	_2M	_3M	_4M	_5M	_6M	_9M	_12M	_18M	_24M
Mean	0.0004	0.1927	0.2044	0.2130	0.2190	0.2238	0.2278	0.2357	0.2412	0.2503	0.2578
Std	0.0116	0.0892	0.0876	0.0850	0.0817	0.0792	0.0769	0.0714	0.0669	0.0627	0.0592
Min	-0.094	0.0924	0.1108	0.1194	0.1274	0.1342	0.1374	0.1482	0.1542	0.1645	0.1709
Max	0.0684	0.7754	0.7445	0.7013	0.6739	0.6504	0.6314	0.5937	0.5518	0.5377	0.5067
Skew	-0.475	2.2261	2.1512	2.0365	1.9125	1.8082	1.7216	1.5453	1.4377	1.2339	1.1198
Kurt	10.647	9.1367	8.5537	7.9602	7.3092	6.7749	6.3677	5.5160	5.0691	4.2227	3.8957

Note. The table presents the mean, standard deviation, minimum, maximum, skewness, and kurtosis for daily index returns and variance swaps (in volatility terms) with maturity of 1–6-, 9-, 12-, 18- and 24-month over the period November 2008–September 2017.

In contrast to the index process, variance swaps exhibit positive skewness and excess kurtosis across the entire maturity spectrum, as presented in Table 2.1. Besides these time-series characteristics, the variance swap rates convey an additional layer of information along the term structure dimension. The mean values of variance swaps show an upward trend with maturity while higher-order statistics exhibit opposite patterns, with shorter-term variance swap displaying greater variability, asymmetry and tailedness. Figure 2.2 corroborates this observation by

indicating that the turbulence causes an inversion of the term structure with the shortest maturity responding most to market turmoil. During periods of relative calm, for example November 2011 to August 2015, the long-end of the term structure is on average 3.03% larger than the short-end; while during episodes of market crashes, for instance in November 2008, the short-end spikes up to an average value of 36.36%, exceeding the long-end by 17.31%.

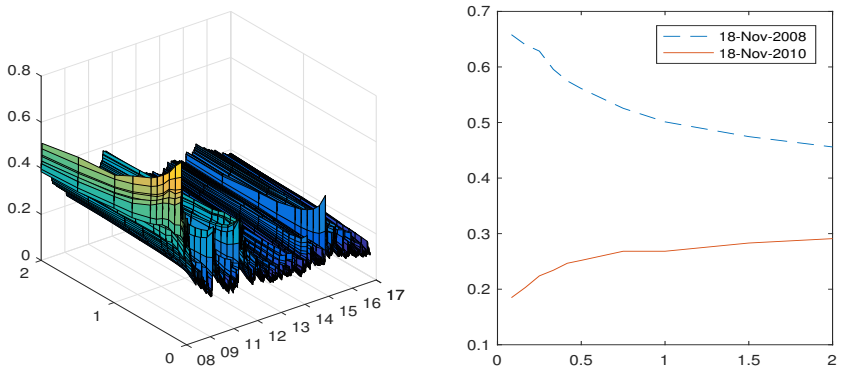

Figure 2.2: The term structure of variance swaps. The figure plots the daily variance swap rates (in volatility terms) on S&P 500 with maturity of 1–6-, 9-, 12-, 18- and 24-month from November 2008 to September 2017 (left panel) and on two specific dates (right panel).

The literature on variance swaps generally summarises the variation in variance swaps across the maturity spectrum in terms of a relatively small number of factors. Egloff, Leippold, and Wu (2010) employ two stochastic variance risk factors to characterise the dynamics of variance swaps with 2-, 3-, 6-, 12-, and 24-month maturities. With an almost identical dataset, Filipović, Gourier, and Mancini (2016) conclude that the first two principal components effectively depict the behaviour characteristics of variance swaps. Despite that two- and three-factor models have achieved satisfactory performance in producing the empirical shapes of variance swap curves, this does not conclusively rule out the possibility that the remaining components contain pertinent information concerning the cross-sectional stock returns. Whether these components correlate with changes in investment opportunities, thus affording valuable insight into expected stock returns, has remained unexplored so far.

To extract the variance swap components, this chapter adopts the principal component approach and selects the optimal number of components that balances the trade-off between accuracy and parsimony, without being restricted to a predetermined number. To achieve this, Table 2.2 compares the predictive performance of different specifications in terms of Akaike's Information Criterion (AIC), which penalises the complexity for adding extra terms. For ease of reference, the following

sections use the notation of VS_i to denote the i-th principal component of variance swaps and VSi to denote the factor set consisting of the first i components.

Table 2.2: Model selection for the value-weighted portfolios

	Fixed-Window					Rolling-Window				
Factors	1M	3M	6M	12M	24M	1M	3M	6M	12M	24M
Mkt	31.05	55.18	75.74	90.11	110.62	27.02	53.88	74.14	94.81	114.97
MktVS1	30.38	54.33	74.37	88.89	109.35	25.98	53.76	74.38	95.26	115.60
MktVS2	28.48	51.27	71.90	88.23	109.39	24.44	52.82	73.15	93.49	113.73
MktVS3	28.83	50.80	71.91	89.03	109.92	23.44	51.99	71.21	92.92	112.57
MktVS4	28.15	50.42	71.14	87.29	106.50	22.48	50.12	70.61	92.62	111.81
MktVS5	22.38	47.36	65.90	79.97	96.62	21.28	48.77	69.70	90.97	110.81
MktVS6	20.84	**46.77**	**65.68**	79.97	**96.53**	19.97	47.92	68.65	90.54	109.97
MktVS7	21.35	47.08	66.04	**79.75**	96.59	19.12	47.57	67.51	90.14	110.80
MktVS8	20.79	46.83	65.86	80.23	97.88	17.87	46.50	66.67	89.70	109.70
MktVS9	20.09	47.44	66.81	81.06	99.20	17.22	45.73	65.82	88.80	108.60
MktVS10	**19.47**	48.03	68.20	82.70	100.81	**15.86**	**44.09**	**63.71**	**86.70**	**106.61**
FF3	16.81	45.06	63.04	80.91	98.39	13.80	42.11	60.76	82.31	102.53
FF3VS1	17.23	45.73	63.64	80.94	98.39	13.96	42.36	61.71	83.29	103.09
FF3VS2	15.05	43.25	61.78	78.16	93.28	13.97	42.62	62.05	83.84	103.28
FF3VS3	15.32	43.57	62.65	79.15	94.69	14.03	42.25	62.25	84.79	103.89
FF3VS4	15.64	43.22	62.03	78.32	**92.19**	14.03	41.67	61.86	83.91	103.23
FF3VS5	15.44	43.15	61.37	**77.66**	92.27	13.69	41.18	61.23	82.95	102.50
FF3VS6	14.90	43.02	61.35	78.34	92.64	13.58	41.17	60.85	82.41	102.85
FF3VS7	15.03	43.48	61.82	78.25	92.43	13.30	41.41	60.95	82.64	103.89
FF3VS8	14.28	**42.87**	**61.17**	78.33	92.40	12.40	41.54	61.34	83.01	103.82
FF3VS9	14.70	43.74	62.30	79.79	94.14	12.06	41.43	60.63	82.07	103.06
FF3VS10	**14.05**	44.08	63.52	81.23	95.69	**10.76**	**39.90**	**59.56**	**81.26**	**101.72**

Note. The table reports the Akaike's information criterion (AIC) values for predicting the cross-sectional returns of the 25 value-weighted size and book-to-market sorted portfolios. The best model (the smallest AIC) for each specification is indicated in bold. The upper/lower panel presents the AIC values for incorporating the variance swap components to the market factor/Fama–French three factors, respectively. For each model, forecasts of 1M-, 3M-, 6M-, 12M-, and 24M-period excess returns using fixed and rolling schemes are examined. The whole sample period is from Nov 2008 to Sep 2017. In the fixed-window analysis, betas are calculated over the period from Nov 2008 to Apr 2013, while forecasts are assessed over the period from May 2013 to Sep 2017. In the rolling-window analysis, betas are sequentially calculated using daily observations over the previous 12 months.

With panels corresponding to forecasting schemes (fixed and rolling) and columns to horizons (1M, 3M, 6M, 12M, 24M), Table 2.2 examines the predictive performance of incorporating variance swap components to the market factor model (Mkt) and Fama–French three-factor model (FF3) in terms of AIC. The table highlights the AIC values that most effectively predict the cross-sectional returns of the 25 value-weighted factor-sorted portfolios[2]. A consistent pattern observed across

[2]Rather than turning to the combined use of the AIC and Bayesian Information Criterion (BIC), this chapter follows previous research that explicitly distinguishes between explanatory and predictive modelling, and favourably selects AIC as the criterion from a predictive viewpoint (Sober,

all scenarios is that the few remaining components, despite contributing little to explaining variation in the variance swap space compared with the first three components, offer valuable information regarding the cross-sectional stock returns. This phenomenon becomes more pronounced with the rolling-window analysis, where the Mkt/FF3 factors combined with all components generate lowest AIC values and thus represent the most preferred set of predictors, regardless of forecast horizons. For the fixed-window scheme, the shorter the forecast horizon, the more prominent role the few remaining components play. Generally speaking, at least five principal components should be incorporated to exploit the predictive value of the term structure of variance swaps, according to the information criterion. In most cases, introducing more than five principal components results in lower information criterion values, which also applies to the equal-weighted factor-sorted portfolios (see Appendix 2.A).

2.4.2 Model Comparison

The Fama–French three-factor model has empirically demonstrated significant cross-sectional explanatory power for the factor-sorted portfolio returns, thus representing a broadly-recognised benchmark for cross-sectional analysis. The Carhart four-factor model (Carhart, 1997) accounts for the momentum-related anomalies through incorporation of momentum (UMD) as an additional factor. This section presents a comparison of the variance swap components and the established asset-pricing factors in terms of cross-sectional stock predictability, including the market factor (Mkt), the momentum factor combined with the market factor (Mom), the Fama–French three (FF3) factors, and the Fama–French–Carhart (FFC) factors. More specifically, Table 2.3 and Table 2.4 investigate the cross-sectional predictability of the 25 value-weighted and equal-weighted factor-sorted portfolios implied by different predictor sets augmented with variance swap components[3]. As R^2 and the root mean square error (RMSE) that measures the average forecasting error for each portfolio generally improve with addition of more variables, predictability is assessed mainly in terms of the adjusted R^2, R^2_{adj}, to penalise this effect.

In line with the existing research, the size, value, and momentum factors contribute significantly to the predictive power on cross-sectional differences in equity returns, while the predictability generally deteriorates with the longer horizon. Comparing the benchmark factors (upper panel) with the augmented factor inputs

2002; Konishi and Kitagawa, 2008; Shmueli et al., 2010; Berk, 2016).

[3]For ease of discussion, Table 2.3 and Table 2.4 report the results of incorporating the first eight principal components in the fixed-window analysis and the entire set of variance swap components in the rolling-window analysis, as suggested by the information criterion in Table 2.2.

(lower panel), Table 2.3 and Table 2.4 indicate that incorporating variance swap information substantially improves R^2_{adj} in all cases. For example, extending the Mkt factor model to include variance swap components as additional predictor variables results in an increase in average R-squared of the 1-month-ahead returns from 0.73% (first column in the upper-left panel) to 44.63% (lower-left panel).

Table 2.3: Predicting the value-weighted size and book-to-market sorted portfolios

	Fixed-Window				Rolling-Window			
	Mkt	Mom	FF3	FFC	Mkt	Mom	FF3	FFC
RMSE_1m	0.2076	0.1547	0.1331	0.1207	0.2077	0.2160	0.1612	0.1556
R^2_{adj}_1m	0.73%	16.94%	43.72%	44.50%	5.11%	16.05%	43.29%	46.41%
_3m	-0.48%	16.61%	32.89%	34.73%	2.95%	15.14%	38.47%	42.01%
_6m	-0.74%	19.80%	40.25%	40.75%	-2.24%	10.32%	39.92%	41.80%
_12m	-3.30%	3.66%	31.10%	32.80%	-11.55%	3.56%	34.35%	36.23%
_24m	-19.28%	-9.23%	27.36%	30.49%	-24.38%	-7.89%	28.29%	34.11%
	Mkt+VS	Mom+VS	FF3+VS	FFC+VS	Mkt+VS	Mom+VS	FF3+VS	FFC+VS
RMSE_1m	0.1433	0.0949	0.0906	0.0892	0.0941	0.0963	0.0926	0.0921
R^2_{adj}_1m	44.63%	49.52%	55.14%	56.17%	47.88%	50.08%	55.96%	57.77%
_3m	38.48%	44.96%	46.49%	49.08%	45.37%	49.29%	51.63%	54.24%
_6m	43.02%	51.41%	51.89%	54.70%	44.62%	48.10%	51.12%	54.24%
_12m	40.81%	46.35%	46.36%	49.25%	34.45%	38.00%	46.20%	46.85%
_24m	40.74%	49.69%	55.27%	57.35%	26.53%	26.70%	39.06%	40.60%

Note. This table investigates the cross-sectional predictability of the 25 value-weighted portfolios implied by established asset pricing factors (market, size, value, and momentum) and their combination with variance swap components. The whole sample period is from Nov 2008 to Sep 2017. In the fixed-window analysis, betas are calculated over the period from Nov 2008 to Apr 2013, while forecasts are assessed over the period from May 2013 to Sep 2017. In the rolling-window method, betas are sequentially calculated using daily observations over the previous 12 months. For each model, the table reports the root mean square error (RMSE), which measures the average forecasting error of 1-month-ahead returns, and the average adjusted R^2 (R^2_{adj}) for future excess returns at all horizons.

The outperformance of the FF3+VS and FFC+VS specifications relative to the FF3 and FFC benchmarks lends support to the predictive value of variance swaps above and beyond those afforded by the Fama–French–Carhart factors, with the improvements in performance observed across horizons (1M, 3M, 6M, 12M, 24M) and forecasting schemes (fixed and rolling). In the case of a fixed window scheme (left panel), the most effective forecasts are achieved by the combination of the established factors and variance swap components, producing the highest R-squared of 56.17% on 1-month future returns, compared with 44.50% for FFC, 43.72% for FF3, 16.94% for Mom, and 0.73% for Mkt. Of particular note in Table 2.3 is that the employment of variance swap components in conjunction with the market factor alone produces comparable performance to the Fama–French–Carhart factors (FFC) across different scenarios, suggesting considerable predictive validity

of the variance swap components.

Collaborating these findings, Figure 2.3 presents a visual comparison among different model specifications, which plots the realised average excess returns for the 25 portfolios against the corresponding predictions produced by different models. The closer these ratios lie to the 45-degree line through the origin, the better the forecast performance. The Mkt+VS model (bottom-left panel) shows considerable improvement over the Mkt model (top-left panel) and exhibits comparable performance to the FF3 model (middle-left panel), scattering closely around the 45-degree line. Moreover, incorporating variance swap elements substantially enhances the prediction of the most extreme quintiles, such as the 11- and 21-portfolios, which feature smaller sizes and lower book-to-market ratios. Being more closely related to companies' earnings, growth stocks have greater susceptibility to uncertainty in terms of economic and business conditions, concerning which the variance swap factors offer additional insights.

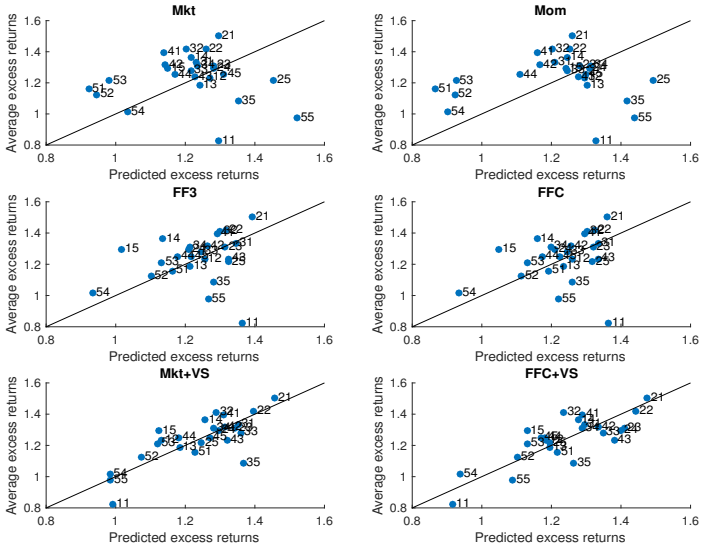

Figure 2.3: Predicting the 25 size and book-to-market sorted portfolios. The realised average excess returns for the 25 portfolios are plotted against the predicted excess returns produced by different specifications. Label (i,j) refers to the portfolio that falls into the i-th smallest size category and j-th lowest book-to-market category.

The performance benefits of variance swap components are even more prominent in the case of the equal-weighted factor-sorted portfolios. Compared to the value-weighted portfolios, the R^2_{adj} of 12-month future returns on the equal-weighted portfolios is 30% lower on average, as seen in Table 2.4. The differences in findings regarding the weighting schemes may be attributed to the relatively higher de-

gree of uncertainty and instability associated with smaller companies. In line with previous conclusions for the value-weighted portfolios (Table 2.3), the combined set of FFC+VS factors demonstrates superior forecasting performance even when compared with Fama–French–Carhart factors, showing an average improvement of 29.65% in R^2_{adj} over fixed windows and 19.70% over rolling windows.

Table 2.4: Predicting the equal-weighted size and book-to-market sorted portfolios

	Fixed-Window				Rolling-Window			
	Mkt	Mom	FF3	FFC	Mkt	Mom	FF3	FFC
RMSE_1m	0.28067	0.24686	0.24538	0.24459	0.28450	0.27910	0.27823	0.25394
R^2_{adj}_1m	-5.49%	11.29%	37.86%	37.98%	-0.63%	8.98%	33.99%	36.58%
_3m	-12.08%	4.07%	23.22%	23.73%	-10.78%	1.31%	23.06%	28.57%
_6m	-16.13%	2.28%	25.26%	24.69%	-21.64%	-3.93%	17.80%	25.07%
_12m	-26.04%	-17.94%	3.59%	1.71%	-41.57%	-22.72%	-2.17%	5.99%
_24m	-18.31%	-9.96%	15.03%	12.38%	-48.33%	-36.54%	-19.19%	-12.00%
	Mkt+VS	Mom+VS	FF3+VS	FFC+VS	Mkt+VS	Mom+VS	FF3+VS	FFC+VS
RMSE_1m	0.18610	0.14421	0.11619	0.11453	0.14636	0.13876	0.14536	0.12960
R^2_{adj}_1m	45.83%	52.67%	55.43%	55.79%	43.20%	45.43%	51.78%	53.56%
_3m	34.70%	44.79%	48.03%	50.25%	34.26%	39.12%	40.94%	44.95%
_6m	39.51%	50.14%	53.04%	55.48%	31.89%	36.77%	39.77%	43.26%
_12m	35.51%	41.30%	46.20%	47.94%	13.22%	17.63%	22.07%	24.30%
_24m	21.49%	27.40%	38.80%	39.28%	5.37%	3.46%	18.40%	16.64%

Note. This table investigates the cross-sectional predictability implied by established asset pricing factors (market, size, value, and momentum) and their combination with variance swap components for the 25 equal-weighted portfolios as in Table 2.3.

Overall, regardless of forecasting schemes, horizons, and weighting methods, the variance swap components demonstrate strong potential for boosting forecasting capabilities. This predictive content is most effectively exploited when the variance swap information is used in conjunction with benchmark Fama–French–Carhart factors. These conclusions also hold true in the 49 industry portfolios, as reported in Appendix 2.C.

2.5 Sources of Predictive Content

This section seeks to investigate sources of predictive content of variance swap components through exploring their relationship with potential equity premium predictors. Despite the lack of consensus on cross-sectional predictors, the literature on market-level predictors provides important clues toward cross-sectional predictive determinants by identifying variables that represent the investment opportunity set (Petkova, 2006; Campbell and Thompson, 2007; Welch and Goyal, 2007; Pettenuzzo

et al., 2014). This section begins by examining the pairwise correlations between variance swap components and a number of commonly employed predictors before proceeding to discuss the contemporaneous and predictive relations using time-series regressions. Besides historical market moments, option-implied risk-neutral moments are also widely considered in this context as they gauge investors' expectations of future market movements, the information content of which is examined in the following subsection.

2.5.1 Relationship with State Variables

Apart from the traditional financial ratios, several studies have attempted to extract relevant information from the bond market. Short-term Treasury yields and the term spread (the difference between the short- and long-term treasury bonds), which correspond to the level and slope of the yield curve, are incorporated to account for the impact of bond market variations on the investment opportunity set. Variables that are informative in terms of the conditional distributions of asset returns, such as the default return spread and default yield spread, are also candidate predictors of the equity premium. Another category of predictors is concerned with the moment-based measures of the market return distributions, including volatility, skewness, and kurtosis. This subsection utilises the sample skewness and kurtosis of historically realised returns, calculated using daily returns over a 6-month rolling window. The calculation of variance follows the approach in Welch and Goyal (2007) by taking the sum of squared daily returns. Option-implied risk-neutral moments are constructed and analysed in the following subsection.

In order to provide a general description of relationships among the factors, Table 2.5 presents the pairwise correlations among variables with the rows corresponding to the variance swap components, VS_1–VS_{10}, and the columns corresponding to the moment measures (Var, Skew, Kurt) and innovations in the state variables, including short-term Treasury-bill yields (u_{tbl}), term spread (u_{tms}), default yield spread (u_{dfy}), default return spread (u_{dfr}), and inflation rate (u_{infl}), where u refers to the innovations in these financial and macroeconomic variables. The upper panel in Table 2.5 examines the contemporaneous relationship while the lower panel takes the lagged values of variance swap components to explore the predictive relations. Overall, Table 2.5 suggests considerable correlations between the variance swap components and almost all variables.

Table 2.5: Pair-wise correlations with state variables

t\t	u_{tbl}	u_{tms}	u_{dfy}	u_{dfr}	u_{infl}	Var	Skew	Kurt
VS_1	**-0.268**	-0.156	0.032	-0.035	0.185	**0.867**	**0.249**	-0.016
VS_2	-0.197	**-0.363**	**0.396**	0.127	0.055	**0.240**	-0.110	**0.212**
VS_3	**-0.385**	**-0.335**	**0.352**	**0.348**	-0.042	**0.321**	0.135	0.000
VS_4	**0.298**	0.059	-0.051	0.122	-0.004	**-0.212**	0.092	-0.134
VS_5	**0.320**	**0.234**	-0.154	-0.084	-0.053	0.036	0.004	0.038
VS_6	0.028	-0.041	0.029	**-0.233**	0.067	0.060	-0.131	0.040
VS_7	0.047	**-0.224**	0.063	0.134	0.090	0.151	**-0.267**	**0.208**
VS_8	**0.218**	-0.040	-0.002	**-0.202**	-0.148	-0.084	**-0.209**	0.112
VS_9	-0.089	-0.038	-0.115	0.098	0.101	**0.246**	0.181	-0.066
VS_{10}	0.161	0.145	-0.159	0.051	0.108	0.149	0.172	-0.102
t\t+1	u_{tbl}	u_{tms}	u_{dfy}	u_{dfr}	u_{infl}	Var	Skew	Kurt
VS_1	0.003	**0.216**	**-0.357**	-0.180	**0.267**	**0.786**	**0.252**	-0.142
VS_2	0.177	-0.041	0.121	-0.100	**0.217**	**0.243**	-0.161	0.182
VS_3	-0.088	-0.022	0.141	0.061	0.154	**0.436**	0.143	-0.023
VS_4	-0.120	-0.146	**0.263**	**0.396**	-0.152	-0.133	0.081	-0.059
VS_5	-0.173	0.143	0.108	0.061	-0.112	0.150	-0.003	0.060
VS_6	0.138	-0.148	-0.103	0.066	0.006	0.017	-0.055	0.055
VS_7	**0.313**	-0.077	0.020	-0.007	0.078	0.002	**-0.233**	0.188
VS_8	-0.087	-0.187	0.020	**0.227**	-0.063	0.046	-0.135	0.165
VS_9	0.002	0.150	-0.079	-0.181	-0.090	**0.289**	0.160	-0.025
VS_{10}	0.062	0.106	0.029	-0.095	-0.095	0.110	0.062	0.000

Note. This table reports the time-series correlations between each variance swap component (in rows) and the prevalent predictor variables, including innovations in short-term bill yield (u_{tbl}), term spread (u_{tms}), default yield spread (u_{dfy}), default return spread (u_{dfr}), inflation rate (u_{infl}), and various moment measures, including variance (Var), skewness (Skew), and kurtosis (Kurt). Innovations are proxied by the first differences of corresponding variables. The upper panel "t\t" reports the contemporaneous correlations while the lower panel "t\t+1" reports the predictive relationships by examining the one-month lagged variance swap components. Correlations with absolute values higher than 0.2 are marked in bold.

To determine whether a statistically significant relationship exists and to further explore how variance swap components enhance forecast accuracy, Table 2.6 performs regression analysis to assess the explanatory and predictive power

$$F_{t/t+h} = c_0 + cVS_t + \epsilon_t, \tag{2.11}$$

where $F_{t/t+h}$ refers to the current/future values of the dependent variables, and VS_t denotes the entire set of variance swap components. The variance swap components show significant contemporaneous associations with all factors in terms of the F-statistic, while retaining their significant predictive value for all variables that capture shifts in credit market conditions and monetary policy (Abhyankar and Gonzalez, 2009; Boons, 2016).

To facilitate direct comparison with benchmarks, with particular focus on the additional information provided by variance swap components beyond that obtained

from benchmark factors, Table 2.7 follows Petkova (2006) to conduct the following analysis as described by Equation (2.12). This table investigates and compares the performance of the Fama–French factors, momentum factor, and variance swap components in characterising the variations in state variables

$$\begin{aligned} u_{t/t+h} &= c_0 + c_{MKT} MKT_t + c_{SMB} SMB_t + c_{HML} HML_t + \epsilon_t, \\ u_{t/t+h} &= c_0 + c_{MKT} MKT_t + c_{UMD} UMD_t + \epsilon_t, \\ u_{t/t+h} &= c_0 + c_{MKT} MKT_t + cVS_t + \epsilon_t, \end{aligned} \tag{2.12}$$

where $u_{t/t+h}$ stands for the current/future innovations in the financial and macroeconomic variables and $MKT_t/SMB_t/HML_t/UMD_t$ refer to the benchmark market/size/value/momentum factors, respectively.

Table 2.6: Contemporaneous and predictive relations with state variables

	Contemporaneous				Predictive			
	R^2	R^2_{adj}	F_stat	F_pval	R^2	R^2_{adj}	F_stat	F_pval
u_{tbl}	35.22%	28.40%	5.165	0.000	23.88%	15.78%	2.948	0.003
u_{tms}	29.81%	22.42%	4.035	0.000	14.45%	5.34%	1.587	0.122
u_{dfy}	30.41%	23.09%	4.152	0.000	34.44%	27.47%	4.938	0.000
u_{dfr}	31.46%	24.24%	4.360	0.000	23.04%	14.85%	2.814	0.004
u_{infl}	16.45%	7.65%	1.870	0.059	16.93%	8.09%	1.916	0.052
Var	85.56%	84.04%	56.31	0.000	67.89%	64.48%	19.88	0.000
Skew	28.48%	20.95%	3.783	0.000	24.85%	16.86%	3.109	0.002
Kurt	15.72%	6.85%	1.772	0.076	19.13%	10.53%	2.224	0.023

Note. This table presents time-series regressions in Equation (2.11) of the commonly-used predictors on the entire set of variance swap components, including the innovations in the short-term T-bill yields (u_{tbl}), term spread (u_{tms}), default yield spread (u_{dfy}), default return spread (u_{dfr}), inflation rate (u_{infl}), and moment measures of aggregate market returns (Var, Skew, Kurt). The left panel captures the contemporaneous relations while the right panel examines the predictive relations using predictive regressions. Each panel reports the R^2, the adjusted R^2 (R^2_{adj}), the F statistic (F_stat), and the p-value (F_pval).

In terms of contemporaneous relationship, all specifications possess high explanatory power for state variables, with the VS factors replicating more of the variation relative to alternative factors and obtaining an R^2_{adj} of 28.63% on short-term Treasury bill innovations compared with 14.96% in the Momentum model and 15.61% in the FF3 model (Table 2.7). Moreover, the VS factors yield additional information on the higher-order moments, which cannot be adequately accounted for by the other two specifications. The predictive regressions reveal more prominent outperformance of VS information set, which produces moderate R^2_{adj} and presents significant predictive power for all variables except the term spread.

Table 2.7: Comparing the predictive and explanatory power for state variables

	FF3			Mom			VS		
t\t	R^2	R^2_{adj}	F_stat	R^2	R^2_{adj}	F_stat	R^2	R^2_{adj}	F_stat
u_{tbl}	18.02%	15.61%	0.000**	16.58%	14.96%	0.000**	36.10%	28.63%	0.000**
u_{tms}	17.79%	15.38%	0.000**	17.69%	16.01%	0.000**	32.52%	24.62%	0.000**
u_{dfy}	20.15%	17.81%	0.000**	22.86%	21.36%	0.000**	31.70%	23.71%	0.000**
u_{dfr}	10.41%	7.77%	0.010**	11.63%	9.92%	0.002**	42.53%	35.80%	0.000**
u_{infl}	6.21%	3.45%	0.087*	7.68%	5.89%	0.016**	20.00%	10.64%	0.025**
Var	24.26%	22.03%	0.000**	17.83%	16.23%	0.000**	86.44%	84.85%	0.000**
Skew	1.66%	-1.23%	0.633	0.07%	-1.87%	0.966	28.92%	20.60%	0.000**
Kurt	5.55%	2.77%	0.119	5.27%	3.43%	0.062*	17.65%	8.01%	0.059*
t\t+1	R^2	R^2_{adj}	F_stat	R^2	R^2_{adj}	F_stat	R^2	R^2_{adj}	F_stat
u_{tbl}	4.19%	1.35%	0.226	2.78%	-0.88%	0.237	25.34%	16.51%	0.003**
u_{tms}	1.52%	-1.41%	0.671	2.54%	0.63%	0.269	14.45%	4.33%	0.174
u_{dfy}	2.15%	-0.76%	0.531	10.53%	8.77%	0.003**	38.10%	30.78%	0.000**
u_{dfr}	2.91%	0.03%	0.392	3.11%	1.21%	0.200	23.20%	14.12%	0.007**
u_{infl}	6.95%	4.19%	0.063*	6.35%	4.52%	0.035**	17.07%	7.26%	0.077*
Var	24.49%	22.24%	0.000**	23.76%	22.26%	0.000**	67.99%	64.21%	0.000**
Skew	0.75%	-2.20%	0.858	0.52%	-1.43%	0.768	26.87%	18.22%	0.001**
Kurt	2.33%	-0.57%	0.495	1.12%	-0.82%	0.564	19.34%	9.80%	0.034**

Note. This table presents the time-series regressions in Equation (2.12) of the state variables, including the innovations in the short-term T-bill yields (u_{tbl}), term spread (u_{tms}), default yield spread (u_{dfy}), default return spread (u_{dfr}), inflation rate (u_{infl}), and the moment measures of aggregate market returns (Var, Skew, Kurt), on different set of factors. The left panel corresponds to the Fama–French three factors (FF3), the middle panel to the momentum factor combined with the market factor (Mom), and the right panel to the variance swap components combined with the market factor (VS). The table examines the explanatory power (upper panel) and the predictive power (lower panel) in terms of R^2, adjusted R^2 (R^2_{adj}), and F statistic (F_stat) with */** indicating significance at 10%/5% level.

2.5.2 Option-Implied Risk-Neutral Moments

This subsection seeks interpretations of the information content of variance swap components in terms of the risk-neutral moments. To retrieve risk-neutral densities from option prices, this subsection adopts the broadly-employed model-free approach proposed by Bakshi, Kapadia, and Madan (2003) (Chang et al., 2013; Conrad et al., 2013; An et al., 2014). Details for constructing the measures of volatility, skewness, and kurtosis from option prices are described in the Appendix 2.D .

As suggested in Table 2.8, the explanatory and predictive power of variance swap components is substantially more pronounced for risk-neutral measures than for physical measures; for example the R^2_{adj} for the kurtosis increases from 6.85% to 63.23% for contemporaneous connections and from 10.53% to 51.75% for predictive associations. While possessing considerable explanatory power for innovations in all moment measures and notable predictive power for the implied volatility inno-

vations u^{Q}_{Var}, the variance swap components barely contain any information about unanticipated skewness and kurtosis shocks u^{Q}_{Skew} and u^{Q}_{Kurt}.

Table 2.8: Explanatory and predictive power for moment measures

	Contemporaneous				Predictive			
	R^2	R^2_{adj}	F_stat	F_pval	R^2	R^2_{adj}	F_stat	F_pval
Var^{P}	85.56%	84.04%	56.31	0.000**	67.89%	64.48%	19.88	0.000**
$Skew^{P}$	28.48%	20.95%	3.783	0.000**	24.85%	16.86%	3.109	0.002**
$Kurt^{P}$	15.72%	6.85%	1.772	0.076*	19.13%	10.53%	2.224	0.023**
Var^{Q}	99.83%	99.81%	5455	0.000**	81.42%	79.44%	41.19	0.000**
$Skew^{Q}$	66.29%	62.74%	18.68	0.000**	59.14%	54.80%	13.61	0.000**
$Kurt^{Q}$	66.73%	63.23%	19.05	0.000**	56.39%	51.75%	12.16	0.000**
u^{P}_{Var}	35.29%	28.48%	5.180	0.000**	36.14%	29.34%	5.319	0.000**
u^{P}_{Skew}	14.82%	5.85%	1.653	0.104	4.11%	-6.09%	0.403	0.942
u^{P}_{Kurt}	22.86%	14.74%	2.815	0.004**	3.73%	-6.52%	0.364	0.959
u^{Q}_{Var}	79.01%	76.80%	35.76	0.000**	30.27%	22.85%	4.081	0.000**
u^{Q}_{Skew}	24.41%	16.46%	3.068	0.002**	12.14%	2.79%	1.299	0.243
u^{Q}_{Kurt}	15.21%	6.28%	1.704	0.091*	7.62%	-2.21%	0.775	0.652

Note. This table presents time-series regressions in Equation (2.11) of the physical (denoted by P) and risk-neutral (denoted by Q) moment measures at 180-day horizon and their innovations (u_{Var}, u_{Skew}, u_{Kurt}) on the entire set of variance swap components.

When combined with the market factor (Table 2.9), the variance swap components display strong associations with a range of moment measures, especially risk-neutral ones, as suggested by the R^2_{adj} of 99.81% for the option-implied volatility and R^2_{adj}s over 60% for the option-implied skewness and kurtosis. Despite being slightly weaker, the predictive relations are still substantial, as indicated by the R^2_{adj} of 79.32% for the option-implied volatility and R^2_{adj}s over 50% for the option-implied skewness and kurtosis. Comparing across models, the variance swap components exhibit higher degrees of co-movement with all moment measures, especially for the risk-neutral metrics, explaining up to 63.07% of the variation in implied skewness compared with 3.08% (FF3) and 2.94% (Mom), and predicting up to 51.34% of the variation in implied kurtosis compared with 5.39% (FF3) and 0.86% (Mom).

Table 2.10 further assesses the predictive and explanatory power for the innovations in moment measures, u_{Var}, u_{Skew}, u_{Kurt}. Despite achieving similar explanatory power for the innovations in option-implied moments, the FF3 and Mom carry negligible information in terms of predicting future innovations (lower panel) in all cases, while the VS components possess considerable predictive power for the future innovations in risk-neutral volatility u^{Q}_{Var}. Future innovations in higher-order moments, whether under physical or risk-neutral measures, cannot be effectively captured by the listed factors.

Table 2.9: Comparing the predictive and explanatory power for moment measures

	FF3			Mom			VS		
t\t	R^2	R^2_{adj}	F_stat	R^2	R^2_{adj}	F_stat	R^2	R^2_{adj}	F_stat
Var^{P}	24.26%	22.04%	0.000**	17.83%	16.24%	0.000**	86.44%	84.86%	0.000**
Skew^{P}	1.66%	-1.23%	0.633	0.07%	-1.87%	0.966	28.92%	20.60%	0.000**
Kurt^{P}	5.55%	2.77%	0.119	5.27%	3.43%	0.062*	17.65%	8.01%	0.059*
Var^{Q}	26.33%	24.16%	0.000**	21.97%	20.46%	0.000**	99.83%	99.81%	0.000**
Skew^{Q}	6.19%	3.43%	0.088*	4.97%	3.13%	0.072*	66.62%	62.72%	0.000**
Kurt^{Q}	5.85%	3.08%	0.103	4.79%	2.94%	0.080*	66.94%	63.07%	0.000**
t\t+1	R^2	R^2_{adj}	F_stat	R^2	R^2_{adj}	F_stat	R^2	R^2_{adj}	F_stat
Var^{P}	24.49%	22.24%	0.000**	23.76%	22.26%	0.000**	67.99%	64.21%	0.000**
Skew^{P}	0.75%	-2.20%	0.858	0.52%	-1.43%	0.768	26.87%	18.22%	0.001**
Kurt^{P}	2.33%	-0.57%	0.495	1.12%	-0.82%	0.564	19.34%	9.80%	0.034**
Var^{Q}	9.65%	6.96%	0.016**	6.24%	4.40%	0.037*	81.51%	79.32%	0.000**
Skew^{Q}	7.38%	4.63%	0.051*	3.13%	1.23%	0.197	59.41%	54.60%	0.000**
Kurt^{Q}	8.12%	5.39%	0.035**	2.76%	0.86%	0.240	56.48%	51.34%	0.000**

Note. This table presents the time-series regressions in Equation (2.12) of the moment measures on three sets of factors. The left panel corresponds to the Fama–French three factors (FF3), the middle panel to the momentum factor combined with the market factor (Mom), and the right panel to the variance swap components combined with the market factor (VS).

Table 2.10: Comparing the predictive and explanatory power for innovations in moment measures

	FF3			Mom			VS		
t\t	R^2	R^2_{adj}	F_stat	R^2	R^2_{adj}	F_stat	R^2	R^2_{adj}	F_stat
$\text{u}^{\text{P}}_{\text{Var}}$	9.08%	6.41%	0.021**	3.73%	1.86%	0.141	35.55%	28.01%	0.000**
$\text{u}^{\text{P}}_{\text{Skew}}$	8.37%	5.68%	0.030**	8.51%	6.74%	0.010**	21.11%	11.87%	0.016**
$\text{u}^{\text{P}}_{\text{Kurt}}$	5.47%	2.69%	0.123	4.00%	2.13%	0.122	26.81%	18.24%	0.001**
$\text{u}^{\text{Q}}_{\text{Var}}$	84.17%	83.71%	0.000**	83.97%	83.66%	0.000**	89.48%	88.25%	0.000**
$\text{u}^{\text{Q}}_{\text{Skew}}$	42.49%	40.80%	0.000**	46.04%	44.99%	0.000**	46.66%	40.42%	0.000**
$\text{u}^{\text{Q}}_{\text{Kurt}}$	27.81%	25.69%	0.000**	33.09%	31.80%	0.000**	37.29%	29.95%	0.000**
t\t+1	R^2	R^2_{adj}	F_stat	R^2	R^2_{adj}	F_stat	R^2	R^2_{adj}	F_stat
$\text{u}^{\text{P}}_{\text{Var}}$	9.45%	6.76%	0.018**	5.99%	4.15%	0.043**	38.95%	31.73%	0.000**
$\text{u}^{\text{P}}_{\text{Skew}}$	2.09%	-0.82%	0.543	0.35%	-1.61%	0.838	4.53%	-6.77%	0.952
$\text{u}^{\text{P}}_{\text{Kurt}}$	2.21%	-0.69%	0.518	0.16%	-1.80%	0.921	3.85%	-7.52%	0.975
$\text{u}^{\text{Q}}_{\text{Var}}$	0.76%	-2.19%	0.857	0.39%	-1.56%	0.820	33.76%	25.92%	0.000**
$\text{u}^{\text{Q}}_{\text{Skew}}$	2.82%	-0.06%	0.406	1.89%	-0.03%	0.377	14.30%	4.17%	0.181
$\text{u}^{\text{Q}}_{\text{Kurt}}$	1.61%	-1.32%	0.649	0.64%	-1.31%	0.722	8.70%	-2.10%	0.634

Note. This table presents the time-series regressions in Equation (2.12) of the innovations in moment measures on three sets of factors. The table examines the explanatory power (upper panel) and the predictive power (lower panel) in terms of R^2, adjusted R^2 (R^2_{adj}), and F statistic (F_stat) with */** indicating significance at 10%/5% level.

As seen in Table 2.11, the variance swap components compare favourably with models that utilise different combinations of realised and option-implied mo-

ments, and their innovations. Due to the limited number of options quotes available, the most extended horizon that permits construction of reliable option-implied moments is 6 months. The variance swaps, however, span from 1 month to 2 years and accordingly reflect investors' expectations further into the future, which might justify their outperformance relative to the option-implied moments, particularly over longer time horizons. Overall, the variance swap components present a useful source of information that contributes towards capturing the variation in underlying economic conditions and aggregate market risks at both short and longer horizons.

Table 2.11: Predicting the factor-sorted portfolios using moment measures

	Value-Weighted					Equal-Weighted				
	RMSE	R^2_{adj}				RMSE	R^2_{adj}			
	1M	1M	6M	12M	24M	1M	1M	6M	12M	24M
VIX	0.208	10.68%	-0.52%	-10.28%	-23.82%	0.279	8.54%	-11.92%	-32.95%	-34.29%
u_{VIX}	0.203	18.12%	9.18%	-0.68%	-14.40%	0.273	11.86%	-12.14%	-32.75%	-40.29%
$Skew^Q$	0.196	14.48%	5.11%	-5.13%	-22.12%	0.257	11.77%	-8.30%	-30.83%	-36.88%
u^Q_{Skew}	0.182	12.73%	3.72%	-3.54%	-18.93%	0.284	6.89%	-19.29%	-38.85%	-47.98%
$Kurt^Q$	0.192	13.87%	5.09%	-5.21%	-22.00%	0.26	9.96%	-9.72%	-31.40%	-34.68%
u^Q_{Kurt}	0.195	11.83%	4.13%	-2.78%	-21.20%	0.280	6.56%	-17.37%	-38.19%	-46.81%
All^Q	0.200	24.45%	13.29%	2.63%	-13.16%	0.244	25.46%	5.94%	-18.36%	-25.89%
u^Q_{All}	0.190	25.15%	16.65%	7.17%	-9.64%	0.249	20.94%	0.55%	-27.40%	-33.97%
$All^Q+u^Q_{All}$	0.168	34.62%	26.91%	17.45%	-2.72%	0.204	35.22%	16.29%	-9.20%	-21.92%
Var^P	0.172	11.70%	7.10%	-6.29%	-18.59%	0.258	8.79%	-8.74%	-28.90%	-41.21%
u_{Var}	0.197	11.99%	6.88%	-2.96%	-13.63%	0.263	8.85%	-3.26%	-29.83%	-33.11%
$Skew^P$	0.165	13.44%	6.87%	-2.12%	-15.78%	0.276	8.83%	-12.75%	-30.51%	-36.72%
u^P_{Skew}	0.211	15.95%	2.36%	-6.55%	-17.52%	0.269	9.42%	-16.24%	-37.94%	-44.67%
$Kurt^P$	0.204	12.74%	5.54%	-9.03%	-21.16%	0.275	7.66%	-12.60%	-32.83%	-35.62%
u^P_{Kurt}	0.202	11.73%	-0.98%	-8.81%	-16.91%	0.278	6.66%	-16.87%	-39.32%	-38.39%
All^P	0.170	25.87%	17.19%	9.13%	-2.54%	0.243	19.69%	0.56%	-14.64%	-23.66%
u^P_{All}	0.179	26.95%	15.06%	3.58%	-0.28%	0.235	21.80%	5.49%	-24.40%	-26.12%
$All^P+u^P_{All}$	0.149	40.89%	30.97%	24.56%	21.25%	0.210	32.26%	17.59%	-4.05%	-9.95%
Mkt	0.208	5.11%	-2.24%	-11.55%	-24.38%	0.285	-0.63%	-21.64%	-41.57%	-48.33%
MktVS1	0.211	10.35%	-0.02%	-10.30%	-23.04%	0.271	8.09%	-11.25%	-34.91%	-37.35%
MktVS3	0.152	23.32%	16.26%	7.43%	-2.39%	0.246	20.19%	2.61%	-16.99%	-21.31%
MktVS5	0.145	33.04%	25.19%	17.90%	9.29%	0.228	29.55%	17.44%	-4.24%	-10.08%
MktVS10	0.09	47.88%	44.62%	34.45%	26.53%	0.146	43.20%	31.89%	13.22%	5.37%

Note. This table investigates the cross-sectional predictability of different combinations of realised and option-implied moment measures in terms of 1M-, 3M-, 6M-, 12M-, 24M-period future excess returns for the 25 value-weighted (left) and equal-weighted (right) portfolios in the rolling-window analysis. For each model, the table reports the root mean square error (RMSE) that measures average forecasting error of 1-month-ahead returns for each portfolio, and the average adjusted R^2 (R^2_{adj}) for future excess returns at all horizons.

Collectively, the variance swap components demonstrate superior explanatory and predictive power for the macroeconomic and financial state variables, and for both the physical and risk-neutral moment measures that characterise the variations in aggregate stock market returns, even when compared with the benchmark factors that capture the size, value and momentum effects. This provides the rationale for the observation that the incorporation of the variance swap compo-

nents contributes to enhanced predictability for the cross section of stock returns, as demonstrated in Section 2.4.

2.6 Conclusion

Providing additional information beyond that afforded by the benchmark factors, variance swap components prove practicable for the forecasting of cross-sectional stock returns. In this chapter, the 25 size and book-to-market sorted portfolios are examined for direct comparison with the benchmark Fama–French–Carhart factors. The information content of variance swaps is linked to state variables that are associated with changes in the investment opportunity set and to statistical measures that characterise market variation. Future research can expand the scope to include an exploration of the information content of variance swap components in the context of other candidate state variables, and to examine the performance of variance swap components in predicting the cross section of returns in other asset classes.

2.7 Appendices

2.A Model Selection for the Equal-Weighted Portfolios

Table 2.12: Model selection for the equal-weighted portfolios

	Fixed-Window					Rolling-Window				
	1M	3M	6M	12M	24M	1M	3M	6M	12M	24M
Mkt	29.73	57.74	79.38	95.81	113.73	25.67	55.46	76.74	99.13	121.48
MktVS1	30.45	58.54	79.94	95.85	113.31	23.91	53.82	75.17	98.78	120.56
MktVS2	24.77	52.06	74.66	94.56	113.25	22.30	52.94	74.63	97.88	119.00
MktVS3	22.55	48.77	70.93	91.63	112.93	21.69	52.15	73.40	97.37	118.60
MktVS4	22.15	48.43	69.91	90.44	111.30	20.71	51.34	72.36	97.16	118.36
MktVS5	17.93	**46.56**	**67.12**	86.30	**107.47**	19.70	50.14	71.13	96.25	118.01
MktVS6	18.29	47.61	68.37	87.91	109.36	18.73	49.75	70.80	95.32	117.14
MktVS7	17.96	47.79	68.23	87.42	110.08	17.63	49.23	69.40	94.19	117.07
MktVS8	18.02	48.05	68.49	86.95	108.25	17.11	48.37	68.91	94.56	117.50
MktVS9	**17.50**	47.96	68.12	**85.47**	109.33	16.36	47.64	68.40	93.80	116.55
MktVS10	17.59	48.33	68.97	86.50	110.48	**15.61**	**46.83**	**67.42**	**92.89**	**116.20**
FF3	16.81	48.36	68.23	87.67	106.06	14.71	46.02	66.95	91.20	117.06
FF3VS1	17.12	48.83	68.44	87.42	105.26	13.72	44.50	65.42	89.94	113.48
FF3VS2	12.99	43.72	64.41	86.14	105.14	12.60	43.70	65.07	89.46	112.37
FF3VS3	12.31	43.21	63.69	84.88	104.63	13.12	43.70	65.32	90.10	112.83
FF3VS4	13.04	43.99	64.40	85.55	106.11	13.24	44.20	64.88	89.89	112.71
FF3VS5	13.14	44.50	64.98	86.07	106.36	13.05	43.91	64.71	90.01	112.80
FF3VS6	13.16	43.78	64.17	86.01	104.54	12.52	43.90	64.03	89.00	112.57
FF3VS7	13.18	42.78	62.30	83.98	105.18	12.11	44.20	64.04	**88.97**	112.76
FF3VS8	12.82	42.13	61.74	82.54	**103.49**	11.87	44.12	64.52	89.58	113.27
FF3VS9	12.22	41.98	**61.47**	**81.23**	103.75	11.41	44.19	64.69	89.53	112.59
FF3VS10	**11.73**	**41.93**	62.29	82.00	105.34	**10.52**	**43.17**	**64.01**	89.21	**112.02**

Note. The table reports the Akaike's information criterion (AIC) values for predicting the cross-sectional returns of the 25 equal-weighted size and book-to-market sorted portfolios. The best model (the smallest AIC) for each specification is indicated in bold. The upper/lower panel presents the AIC values for incorporating the variance swap components to the market factor/Fama–French three factors, respectively. For each model, forecasts of 1M-, 3M-, 6M-, 12M-, and 24M-period excess returns using fixed and rolling schemes are examined. The whole sample period is from Nov 2008 to Sep 2017. In the fixed-window analysis, betas are calculated over the period from Nov 2008 to Apr 2013, while forecasts are assessed over the period from May 2013 to Sep 2017. In the rolling-window analysis, betas are sequentially calculated using daily observations over the previous 12 months.

2.B Predicting the 25 Portfolios from 2010 onwards

This section revisits the main analysis in Table 2.3 and Table 2.4 using a sample that excludes the recession and crisis period, while retaining as many observations as possible. Similar conclusions as those indicated in Table 2.3 and Table 2.4 can be drawn from Table 2.13 and Table 2.14.

Table 2.13: Predicting the value-weighted size and book-to-market sorted portfolios

	Fixed-Window				Rolling-Window			
	Mkt	Mom	FF3	FFC	Mkt	Mom	FF3	FFC
RMSE_1m	0.1760	0.1666	0.1300	0.1296	0.1950	0.2232	0.1713	0.1626
R^2_{adj}_1m	5.26%	8.39%	45.88%	48.43%	7.60%	18.02%	42.80%	46.20%
_3m	3.46%	6.72%	34.65%	38.27%	3.16%	14.51%	35.30%	39.34%
_6m	4.79%	8.38%	41.73%	48.12%	3.40%	12.68%	38.99%	41.07%
_12m	-3.80%	-3.41%	31.13%	40.00%	-2.26%	5.66%	31.26%	33.50%
_24m	-16.80%	-17.19%	26.70%	28.85%	-16.21%	-3.24%	24.82%	32.40%
	Mkt+VS	Mom+VS	FF3+VS	FFC+VS	Mkt+VS	Mom+VS	FF3+VS	FFC+VS
RMSE_1m	0.1188	0.1170	0.0985	0.0962	0.0968	0.1045	0.0897	0.0895
R^2_{adj}_1m	33.23%	35.82%	56.41%	57.41%	48.69%	51.28%	56.61%	58.47%
_3m	29.75%	32.41%	44.59%	46.01%	43.46%	47.84%	49.88%	53.11%
_6m	41.15%	42.92%	53.41%	56.72%	44.51%	48.76%	51.63%	55.48%
_12m	31.77%	29.75%	44.95%	50.26%	31.40%	34.75%	43.03%	44.19%
_24m	23.79%	20.29%	37.04%	43.68%	25.06%	24.34%	36.43%	38.95%

Note. This table investigates the cross-sectional predictability of the 25 value-weighted portfolios implied by established asset pricing factors (market, size, value, and momentum) and their combination with variance swap components. The whole sample period is from Jan 2010 to Sep 2017. In the fixed-window analysis, betas are calculated over the period from Jan 2010 to Apr 2013, while forecasts are assessed over the period from May 2013 to Sep 2017. In the rolling-window method, betas are sequentially calculated using daily observations over the previous 12 months. For each model, the table reports the root mean square error (RMSE), which measures the average forecasting error of 1-month-ahead returns, and the average adjusted R^2 (R^2_{adj}) for future excess returns at all horizons.

Table 2.14: Predicting the equal-weighted size and book-to-market sorted portfolios

	Fixed-Window				Rolling-Window			
	Mkt	Mom	FF3	FFC	Mkt	Mom	FF3	FFC
RMSE_1m	0.2519	0.2349	0.2399	0.2301	0.2587	0.2608	0.2628	0.2371
R^2_{adj}_1m	1.81%	4.02%	40.64%	43.12%	1.88%	11.92%	34.88%	37.72%
_3m	-5.14%	-2.27%	26.68%	30.83%	-7.61%	2.73%	21.81%	27.76%
_6m	-9.31%	-6.72%	28.34%	35.66%	-13.97%	-0.37%	17.89%	25.81%
_12m	-22.83%	-21.70%	6.15%	16.05%	-30.68%	-17.79%	-6.00%	2.89%
_24m	-15.13%	-14.84%	13.04%	15.27%	-41.65%	-32.60%	-20.95%	-12.14%
	Mkt+VS	Mom+VS	FF3+VS	FFC+VS	Mkt+VS	Mom+VS	FF3+VS	FFC+VS
RMSE_1m	0.1693	0.1688	0.1585	0.1530	0.1493	0.1396	0.1546	0.1331
R^2_{adj}_1m	39.98%	40.42%	57.29%	57.49%	42.78%	44.91%	50.80%	52.59%
_3m	35.55%	34.82%	47.82%	48.67%	31.52%	36.65%	38.55%	43.01%
_6m	40.76%	39.96%	53.09%	54.80%	30.27%	35.21%	39.14%	42.37%
_12m	25.77%	23.45%	39.43%	43.30%	9.62%	13.79%	16.77%	19.29%
_24m	14.32%	8.91%	29.95%	32.52%	7.89%	6.08%	17.64%	16.00%

Note. This table investigates the cross-sectional predictability of the 25 equal-weighted portfolios implied by established asset pricing factors (market, size, value, and momentum) and their combination with variance swap components. The whole sample period is from Jan 2010 to Sep 2017. In the fixed-window analysis, betas are calculated over the period from Jan 2010 to Apr 2013, while forecasts are assessed over the period from May 2013 to Sep 2017. In the rolling-window method, betas are sequentially calculated using daily observations over the previous 12 months.

2.C Predicting the 49 Industry Portfolios

The improvement over the benchmark Fama–French factors induced by the variance swap components is more prominent in the 49 industry portfolios than in the 25 factor-sorted portfolios, since the latter by construction is associated with the size and value factors. Examining the performance over rolling windows, the exploitation of variance swap information increases the R^2_{adj} values by more than 20% on average. In contrast to the factor-sorted portfolios, the FF3 factors alone cannot obtain satisfactory prediction results for industry portfolios (see the 4.74% in 24-month ahead forecasts for the value-weighted portfolios and -0.08% for the equal-weighted portfolios), leading to more remarkable enhancement of variance swap components (31.35% for the value-weighted portfolios and 22.76% for the equal-weighted portfolios).

Table 2.15: Predicting the cross section of the 49 industry portfolios

	Value-weighted					Equal-weighted				
	RMSE	R^2_{adj}				RMSE	R^2_{adj}			
	1M	1M	6M	12M	24M	1M	1M	6M	12M	24M
Mkt	0.65460	5.17%	3.93%	-1.44%	-17.19%	0.84278	4.00%	0.18%	-8.28%	-18.42%
MktVS3	0.51490	20.03%	22.34%	16.44%	6.79%	0.68204	19.50%	17.62%	7.63%	-0.10%
MktVS5	0.44756	26.27%	27.41%	21.01%	10.25%	0.57360	25.83%	23.95%	15.93%	5.11%
MktVS10	0.32034	36.38%	38.37%	37.03%	27.01%	0.39534	35.50%	35.11%	31.01%	17.38%
Mom	0.46781	18.66%	18.66%	14.22%	4.81%	0.58151	16.67%	14.82%	9.99%	5.70%
MomVS3	0.38249	28.04%	29.94%	25.69%	21.20%	0.48367	28.10%	27.83%	20.57%	15.98%
MomVS5	0.37615	32.45%	32.20%	28.17%	22.75%	0.44950	31.90%	30.51%	24.78%	18.58%
MomVS10	0.28693	39.91%	41.85%	40.85%	35.47%	0.35839	39.71%	38.69%	35.51%	26.25%
FF3	0.50248	18.21%	20.61%	18.50%	4.74%	0.65052	15.19%	14.47%	9.12%	-0.08%
FF3VS3	0.41446	28.54%	31.65%	27.99%	20.59%	0.54187	26.37%	26.46%	19.00%	10.13%
FF3VS5	0.36988	33.69%	35.36%	31.22%	22.47%	0.47791	31.13%	30.60%	24.55%	13.85%
FF3VS10	0.28750	41.20%	43.13%	43.29%	36.09%	0.37937	38.70%	38.96%	35.34%	22.68%
FF4	0.37810	25.62%	29.33%	28.60%	20.52%	0.50009	21.77%	21.04%	18.72%	15.56%
FF4VS3	0.33715	33.82%	37.25%	34.99%	30.17%	0.42911	32.23%	33.30%	28.81%	23.44%
FF4VS5	0.32299	37.51%	38.98%	37.13%	31.66%	0.39791	35.29%	35.11%	32.35%	25.62%
FF4VS10	0.26178	43.74%	45.91%	46.25%	41.92%	0.33669	42.52%	41.84%	39.39%	30.19%

Note. This table investigates the cross-sectional predictability of the 49 value-weighted (left) and equal-weighted (right) portfolios implied by established asset pricing factors (market, size, value, and momentum) and their combination with variance swap components. In the rolling-window method, betas are sequentially calculated using daily observations over the previous 12 months. For each model, the table reports the root mean square error (RMSE), which measures the average forecasting error of 1-month-ahead returns, and the average adjusted R2 (R^2_{adj}) for future excess returns at all horizons.

2.D Option-Implied Risk-Neutral Moments

Analogous to their physical counterparts, the risk-neutral moments are defined as measures of variability under the risk-neutral pricing measure Q

$$\text{VAR}(t,\tau) = E_t^Q[(R(t,\tau) - E_t^Q[R(t,\tau)])^2], \tag{2.13}$$

$$\text{SKEW}(t,\tau) = \frac{E_t^Q[(R(t,\tau) - E_t^Q[R(t,\tau)])^3]}{(E_t^Q[(R(t,\tau) - E_t^Q[R(t,\tau)])^2])^{3/2}}, \tag{2.14}$$

$$\text{KURT}(t,\tau) = \frac{E_t^Q[(R(t,\tau) - E_t^Q[R(t,\tau)])^4]}{(E_t^Q[(R(t,\tau) - E_t^Q[R(t,\tau)])^2])^2}. \tag{2.15}$$

The Bakshi, Kapadia, and Madan (2003) approach draws upon the insight that the risk-neutral measures can be expanded into functions of $E^Q[R(t,\tau)]$, $E^Q[R(t,\tau)^2]$, $E^Q[R(t,\tau)^3]$, $E^Q[R(t,\tau)^4]$, and thus can be expressed using a set of hypothetical option prices that correspond to the second, third, and fourth powers of future log price returns, that is the quadratic, cubic and quartic contracts. Organising the above equations provides the τ-period risk-neutral variance, risk-neutral skewness, and risk-neutral kurtosis

$$\text{VAR}(t,\tau) = e^{r\tau}V(t,\tau) - \mu(t,\tau)^2, \tag{2.16}$$

$$\text{SKEW}(t,\tau) = \frac{e^{r\tau}W(t,\tau) - 3\mu(t,\tau)e^{r\tau}V(t,\tau) + 2\mu(t,\tau)^3}{[e^{r\tau}V(t,\tau) - \mu(t,\tau)^2]^{3/2}}, \tag{2.17}$$

$$\text{KURT}(t,\tau) = \frac{e^{r\tau}X(t,\tau) - 4\mu(t,\tau)e^{r\tau}W(t,\tau) + 6e^{r\tau}\mu(t,\tau)^2V(t,\tau) - 3\mu(t,\tau)^4}{[e^{r\tau}V(t,\tau) - \mu(t,\tau)^2]^2}, \tag{2.18}$$

where

$$\mu(t,\tau) = E_t^Q[R(t,\tau)] = e^{r\tau} - 1 - \frac{e^{r\tau}}{2}V(t,\tau) - \frac{e^{r\tau}}{6}W(t,\tau) - \frac{e^{r\tau}}{24}X(t,\tau). \tag{2.19}$$

$V(t,\tau)$, $W(t,\tau)$, and $X(t,\tau)$ denote the prices of the hypothetical volatility, cubic, and quartic contracts, respectively, which are constructed from second-, third-, and fourth-order payoffs $H[S]$

$$\begin{cases} \text{Volatility contract :} & H[S] = R(t,\tau)^2, V(t,\tau) = E_t^Q[e^{-r\tau}R(t,\tau)^2], \\ \text{Cubic contract :} & H[S] = R(t,\tau)^3, W(t,\tau) = E_t^Q[e^{-r\tau}R(t,\tau)^3], \\ \text{Quartic contract :} & H[S] = R(t,\tau)^4, X(t,\tau) = E_t^Q[e^{-r\tau}R(t,\tau)^4]. \end{cases}$$

Bakshi and Madan (2000) demonstrate that any twice-continuously differentiable payoff function can be replicated by employing a continuum of out-of-money Euro-

pean calls and puts. Accordingly, the no-arbitrage prices of the volatility, cubic, and quartic contracts can be recovered by utilising the risk-free bonds, the underlying asset, and a combination of out-of-money calls and puts

$$E_t^Q[e^{-r\tau}H[S]] = (H[\bar{S}] - \bar{S}H_S[\bar{S}])e^{-r\tau} + H_S[\bar{S}]S(t) \tag{2.20}$$

$$+ \int_{\bar{S}}^{\infty} H_{SS}[K]C(t,\tau;K)dK + \int_0^{\bar{S}} H_{SS}[K]P(t,\tau;K)dK. \tag{2.21}$$

Through substituting $H_{SS}[K]$, the second-order derivative of the payoff $H[S]$ with respect to the stock price evaluated at the strike price yields the prices of the volatility, cubic, and quartic contracts

$$V(t,\tau) = \int_{S(t)}^{\infty} \frac{2(1-\ln[\frac{K}{S(t)}])}{K^2} C(t,\tau;K)dK + \int_0^{S(t)} \frac{2(1+\ln[\frac{S(t)}{K}])}{K^2} P(t,\tau;K)dK, \tag{2.22}$$

$$W(t,\tau) = \int_{S(t)}^{\infty} \frac{6\ln[\frac{K}{S(t)}] - 3(\ln[\frac{K}{S(t)}])^2}{K^2} C(t,\tau;K)dK - \int_0^{S(t)} \frac{6\ln[\frac{S(t)}{K}] + 3(\ln[\frac{S(t)}{K}])^2}{K^2} P(t,\tau;K)dK, \tag{2.23}$$

$$X(t,\tau) = \int_{S(t)}^{\infty} \frac{12(\ln[\frac{K}{S(t)}])^2 - 4(\ln[\frac{K}{S(t)}])^3}{K^2} C(t,\tau;K)dK + \int_0^{S(t)} \frac{12(\ln[\frac{S(t)}{K}])^2 + 4(\ln[\frac{S(t)}{K}])^3}{K^2} P(t,\tau;K)dK. \tag{2.24}$$

The historical S&P 500 option data $C(t,\tau,K)$ and $P(t,\tau,K)$ are collected from OptionMetrics, available through Wharton Research Data Services. Options with zeros bids, calls with strikes lower than 50% and puts with strikes higher than 150% are filtered out. In-the-money puts and calls are removed due to liquidity concerns, that is, only options with deltas ranging from -0.5 to 0.5 are retained. For each date, a wide variety of options are available at different maturities and moneyness levels. Take 29/09/2017 for example, the maturity of approximately 8,500 options varies from the shortest of 3 days to the longest of around 2 years, while the moneyness ranges from 3.97% to 138.92%. Since not every observation point affords such an adequate number of option quotes, this chapter only considers dates and maturities at which at least four moneyness points are available, which focuses the analysis at the 30-, 90-, and 180-day horizons.

To handle discrete observations, the cubic Hermite spline is employed to obtain a continuum of implied volatility across moneyness levels. Meanwhile, mon-

eyness levels beyond the range of the available data are filled with the lowest or highest implied volatility. Such interpolation/extrapolation procedures are carried out on a thousandths grid that covers moneyness levels (K/S) from 33% to 300% on each date and for each maturity. The filtered implied volatilities are then transformed into call and put prices, moneyness levels greater than 100% for OTM calls and less than 100% for OTM puts, based on which trapezoidal numerical integration is performed to obtain the prices of hypothetical contracts. Moments at fixed 30-, 90-, and 180-day horizons are computed by applying linear interpolation on the available maturities.

Table 2.16 provides a brief description of the option-implied moments, indicating that the option-implied return distributions exhibit substantial negative skewness and excess kurtosis across the spectrum. The implied volatility increases as the horizon expands, while the distribution of volatility itself shows strong positive skewness and kurtosis, as suggested by the quantile values. Consistent with this observation, the skewness measures for the implied volatility are 4.10, 3.46, 3.38, and the kurtosis measures are 25.8, 18.7, 25.4, implying long right-tailed distributions.

Table 2.16: Summary statistics of the option-implied risk-neutral moments

Horizon	Statistics	Var^{Q}	Skew^{Q}	Kurt^{Q}	Observation
30	Mean	5.270%	-2.839	24.97	116
	Median	3.007%	-2.677	20.75	113
	5%	1.325%	-4.692	8.781	49
	95%	18.59%	-1.594	57.93	175
90	Mean	5.637%	-2.491	15.77	94
	Median	3.675%	-2.406	14.35	89
	5%	1.870%	-3.583	6.437	34
	95%	16.89%	-1.549	28.80	164
180	Mean	6.026%	-2.173	10.92	44
	Median	4.434%	-2.183	10.737	46
	5%	2.333%	-2.929	4.374	23
	95%	15.94%	-1.162	17.44	61

Note. This table reports the summary statistics (the mean, the median, the 5% and 95% quantiles) of the option-implied variance, skewness, and kurtosis (Var^{Q}, Skew^{Q}, Kurt^{Q}) at fixed 30-, 90-, and 180-day horizons (in upper, middle, and lower panels) over the sample period from November 2008 to September 2017. The average number of option observations used to recover the moment measures for each date and each horizon is also reported.

Table 2.17: Pair-wise correlations between variance swap components and moment measures

t\t	Var^P	$Skew^P$	$Kurt^P$	Var^Q_{1M}	Var^Q_{3M}	Var^Q_{6M}	$Skew^Q_{1M}$	$Skew^Q_{3M}$	$Skew^Q_{6M}$	$Kurt^Q_{1M}$	$Kurt^Q_{3M}$	$Kurt^Q_{6M}$
VS_1	**0.867**	**0.249**	-0.016	**0.978**	**0.996**	**0.998**	**0.510**	**0.664**	**0.768**	**-0.428**	**-0.620**	**-0.737**
VS_2	**0.240**	-0.110	**0.212**	**0.272**	0.160	0.042	**-0.242**	**-0.206**	-0.174	**0.232**	**0.258**	**0.263**
VS_3	**0.321**	0.135	0.000	**0.484**	**0.444**	**0.427**	**0.416**	**0.306**	**0.291**	**-0.400**	**-0.329**	**-0.329**
VS_4	**-0.212**	0.092	-0.134	**-0.212**	**-0.225**	**-0.238**	-0.162	**-0.215**	-0.196	0.097	0.185	0.195
VS_5	0.036	0.004	0.038	0.106	0.111	0.110	0.011	0.066	0.126	0.000	-0.035	-0.090
VS_6	0.060	-0.131	0.040	0.090	0.068	0.054	0.036	-0.006	-0.012	-0.026	0.002	-0.005
VS_7	0.151	**-0.267**	**0.208**	-0.033	-0.025	-0.026	**-0.200**	-0.124	-0.091	0.193	0.129	0.095
VS_8	-0.084	**-0.209**	0.112	0.044	0.028	0.014	**-0.227**	-0.123	-0.074	**0.213**	0.121	0.076
VS_9	**0.246**	0.181	-0.066	**0.220**	**0.221**	**0.217**	0.002	0.147	0.189	0.011	-0.085	-0.119
VS_10	0.149	0.172	-0.102	0.117	0.117	0.115	0.052	0.129	0.150	-0.054	-0.102	-0.115
t\t+1	Var^P	$Skew^P$	$Kurt^P$	Var^Q_{1M}	Var^Q_{3M}	Var^Q_{6M}	$Skew^Q_{1M}$	$Skew^Q_{3M}$	$Skew^Q_{6M}$	$Kurt^Q_{1M}$	$Kurt^Q_{3M}$	$Kurt^Q_{6M}$
VS_1	**0.786**	**0.252**	-0.142	**0.773**	**0.815**	**0.839**	**0.408**	**0.595**	**0.710**	**-0.346**	**-0.547**	**-0.672**
VS_2	**0.243**	-0.161	0.182	-0.072	-0.111	-0.175	**-0.300**	**-0.234**	**-0.175**	**0.273**	**0.270**	**0.250**
VS_3	**0.436**	0.143	-0.023	**0.227**	**0.263**	**0.286**	**0.346**	**0.276**	**0.291**	**-0.322**	**-0.271**	**-0.308**
VS_4	-0.133	0.081	-0.059	0.002	-0.025	-0.051	0.069	-0.056	-0.095	-0.077	0.051	0.101
VS_5	0.150	-0.003	0.060	**0.362**	**0.323**	**0.279**	0.058	0.125	0.171	-0.010	-0.075	-0.112
VS_6	0.017	-0.055	0.055	0.019	-0.024	-0.045	-0.012	-0.073	-0.043	0.000	0.061	0.029
VS_7	0.002	**-0.233**	0.188	0.005	0.014	0.012	**-0.239**	-0.091	-0.069	**0.237**	0.084	0.073
VS_8	0.046	-0.135	0.165	0.148	0.130	0.124	-0.078	-0.005	0.002	0.089	0.004	-0.005
VS_9	**0.289**	0.160	-0.025	**0.277**	**0.269**	**0.249**	0.062	**0.203**	**0.257**	-0.035	-0.126	-0.170
VS_10	0.110	0.062	0.000	**0.244**	**0.214**	0.178	0.092	0.164	0.186	-0.069	-0.113	-0.121

Note. This table reports the time-series correlations between variance swap components (in rows) and the moment measures, including the realised moment measures Var^P, $Skew^P$, $Kurt^P$, and option-implied measures Var^Q, $Skew^Q$,$Kurt^Q$. The upper panel "t\t" reports the contemporaneous correlations while the lower panel "t\t+1" reports the predictive relationships by examining the one-month lagged variance swap components. Correlations with absolute values greater than 0.2 are marked in bold.

Chapter 3

Predictability and Stability of Macro-Yields Modelling in the Post-Global Financial Crisis Data[1]

3.1 Introduction

Macro-finance term structure models, extending the standard approaches of term structure modelling to incorporate macroeconomic variables, have recently become a benchmark for jointly modelling the yield curve and macroeconomic dynamics. Despite the broad acknowledgement that the term structure and macroeconomic fundamentals both contain important implications for one another, the mechanisms underlying their interactions and the scope for enhanced predictive content have not been fully explored. This chapter examines and quantifies changes in predictive power of macroeconomic variables for government bond yields in the aftermath of the global financial crisis, and documents a substantial deterioration in yield predictability.

The model proposed by Coroneo, Giannone, and Modugno (2016) incorporates a broad range of macroeconomic variables into term structure modelling and accounts for the "spanning puzzle". That is, a significant portion of the information, although bearing minimal relationships with the cross section of current yields, plays an important role in determining future yields. This macro-yields model provides superior explanatory insight and predictive power over the conventional term structure model and other existing macro-yields approaches over the 1970:M1–2008:M12 period. This chapter examines the findings of Coroneo, Giannone, and Modugno (2016) in an extended set of observations that includes the post-global financial crisis period, spanning from 1970:M1 to 2016:M12.

Despite favouring the incorporation of macroeconomic variables, the in-sample analysis suggests that the model specification of two unspanned macro factors, com-

[1]This chapter was written in 2017, circulated and then submitted to the Journal of Applied Econometrics in 2018.

bined with three Nelson–Siegel yield factors to constitute the five-factor model, is not as distinguishable from the other multi-factor models over this prolonged period as for the pre-2009 data. This can largely be attributed to an across-the-board increase in correlations among macroeconomic variables during economic recession and recovery. The vast majority of common variation across the panel is plausibly captured by the single factor that corresponds to this.

The in-sample and out-of-sample predictive power of the macro-yields model deteriorates considerably in the extended sample, resulting in the specification no longer exhibiting superior performance over the only-yields model and other macro-yields approaches. The decline in predictability is marked by an average 0.37 increase in out-of-sample rMSFE with respect to random walk across the full spectrum. Regarding the excess bond returns, the in-sample R^2 decreases by 15% across maturities while the out-of-sample rMSFE with respect to the expectations hypothesis increases by 0.55 on average. Further analysis finds that, this is primarily due to the excessive downward pressure on yield factors exerted by macro components through the VAR model during the recession, when the macro factors experience dramatic changes in response to the substantial decline in general economic activity. The subsequent strong reversal in the macroeconomic panel propagates through the VAR system to drive up the yield factors, further contributing to forecasting errors. This chapter also considers the use of real-time data and assesses the impact of revisions, finding that the conclusions are robust across data revisions. Caruso and Coroneo (2019) confirm these findings, and further examine the role of interest rate survey expectations in this context.

Moreover, this chapter investigates the causes of the deterioration and the stability of yield dynamics and macro-financial linkages in the context of the macro-finance DFM representation. Through the conditional forecasting approach, selected bond yields across the maturity spectrum are reconstructed after 2008:M10 by projections of the other observables based on pre-crisis dynamics and interdependence. While the in-sample counterfactual analysis reveals wide departure of the model-implied path from the actual evolvement, the out-of-sample counterfactual assessment generates large forecast errors, casting doubt on the validity of the underlying predictive mechanisms. That is, the macro-finance interactions and yield dynamics well characterised by the linear-Gaussian DFM framework in pre-2009 data may not adequately capture the post-crisis features.

These conditional forecasts and scenario analysis further suggest that the incapability of the macro-yields model is not fully attributable to the constrained dynamics at the front-end of the curve, but rather to the more general instability of yield dynamics and macro-finance interdependence. Combining the predictability

and stability analysis by means of both unconditional and conditional forecasting evaluation, this chapter highlights the necessity of accounting for the greater complexity in yield distributions and macro-yields interdependence in macro-finance term structure modelling in the post-crisis period.

The remainder of this chapter is structured as follows. After Section 3.2 reviews the related literature, Section 3.3 discusses the model specification and estimation procedures of the macro-finance term structure modelling, with particular focus placed on the macro-yields model with unspanning restrictions. Meanwhile, Section 3.4 presents the in-sample and out-of-sample prediction results over a sample that extends to the post-global financial crisis period, and further explores the underlying reasons by decomposing the realised forecast errors and conducting impulse response analysis. Finally, Section 3.5 investigates the stability of the implied yield dynamics and macro-yields linkages by means of conditional forecasts and scenario analysis, after controlling for the constrained dynamics of the short-term yields. Details of the datasources and the results over the original timespan are in Appendices.

3.2 Related Literature

Numerous empirical studies on term structure have attempted to explore the joint modelling of the Treasury yield curve and the macroeconomy, and to investigate the mechanisms underlying the macro-financial linkages. Taking macroeconomic information into account leads to considerable improvements in forecasts of the yield curve and excess bond returns. Traditional reduced-form term structure modelling broadly falls into three categories: the class of affine term structure models derived from no-arbitrage arguments, the dynamic Nelson–Siegel approach that enables flexible but parsimonious representation of empirically observed curve shapes, and the affine arbitrage-free Nelson–Siegel model that seeks compatibility between the two. Drawing upon the former two strands of research, macro-finance term structure models establish macro-yield linkages using various factor loading specifications, by imposing no-arbitrage restrictions (Ang and Piazzesi, 2003; Ang et al., 2007; Mönch, 2008; Bikbov and Chernov, 2010; Ang et al., 2011; Joslin et al., 2013, 2014) or utilising curve-fitting techniques (Diebold et al., 2006; Bianchi et al., 2009; Koopman et al., 2010; Exterkate et al., 2013; Van Dijk et al., 2014; Altavilla et al., 2017). A third strand of the term structure research suggests an association between the arbitrage-free and Nelson–Siegel structures (Coroneo et al., 2011; Krippner, 2015) and seeks a compatible representation (Christensen et al., 2009, 2010, 2011). Barely any attention has been paid to the macro-finance modelling in this context, except

the single work by Li, Niu, and Zeng (2011) that directly incorporates observed macroeconomic variables into the arbitrage-free Nelson–Siegel model as factors.

Ruling out riskless arbitrage opportunities across yields of different maturities lays the theoretical foundation for the no-arbitrage pricing of bonds. The earliest work of macro-finance term structure models can be traced back to the work by Ang and Piazzesi (2003), which introduces macro variables into the pricing kernel that determines the cross section of bond returns. Pertaining to theoretical rigour and soundness, the no-arbitrage asset pricing framework enjoys considerable popularity in leading research that attempts to gain insights into bond risk premium (Wright, 2011; Joslin et al., 2014). This is in part because the absence of arbitrage corresponds to the existence of risk-neutral measures, under which the bond risk premium can be explicitly identified. The role of the no-arbitrage specification on forecasting, however, is somewhat controversial, and has not been well addressed. While Favero, Niu, and Sala (2012) observe improved predictions at the medium- and long-end of the yield curve during the 1993:M1–2003:M9 evaluation period, Duffee (2011a) finds that a random-walk level factor outperforms the no-arbitrage and other cross-sectional restrictions.

The other important strand of term structure modelling, the dynamic Nelson–Siegel approach, has been widely acknowledged for its goodness-of-fit and rapid implementation. During the past decade, there has been renewed interest in this method due to its flexibility in accommodating time-varying parameters, enabling further investigation into parameter uncertainty. Bianchi, Mumtaz, and Surico (2009) examine the dynamic evolution of yields utilising a time-varying factor-augmented VAR consisting of three Nelson–Siegel factors and three macroeconomic variables; Koopman, Mallee, and Van der Wel (2010) consider a time-varying loading parameter λ_t as a fourth latent factor and induce a common generalised autoregressive conditional heteroskedasticity volatility component to account for time-varying volatility; while Van Dijk, Koopman, Van der Wel, and Wright (2014) propose a "dynamic shifting-endpoint" Nelson–Siegel model that enables changes in the unconditional mean to accommodate the slowly varying trend. Model uncertainty is another challenging topic in the context of yield curve forecasting. Regarding this, De Pooter, Ravazzolo, and Van Dijk (2010) suggest combining individual forecasts to mitigate uncertainty, while Altavilla, Giacomini, and Ragusa (2017) adopt the dynamic model averaging approach that utilises models of different dimensions.

Macro-finance term structure modelling has achieved fresh prominence, with many arguing for the existence of unspanned factors, suggesting that factors spanning the cross section of current yields do not precisely match those determining future yields. The invertibility assumption implied in the prevailing macro-finance

term structure models has been challenged by a growing body of literature (Wright, 2011; Chernov and Mueller, 2012; Bauer et al., 2014; Joslin et al., 2014; Coroneo et al., 2016). Studies such as that conducted by Duffee (2011b) have demonstrated that certain "hidden" components, which are orthogonal to current yields, are predictive of future short rates and bond returns. Accordingly, macro-finance approaches featuring unspanning properties have been proposed to account for this phenomenon, such as the model in Joslin, Priebsch, and Singleton (2014), which extends the Kim–Wright formulation (Kim and Wright, 2005), and the model in Coroneo, Giannone, and Modugno (2016), which builds upon the dynamic Nelson–Siegel representation. Recently, literature has emerged that reveals contradictory findings regarding this unspanning argument; for example, Bauer and Rudebusch (2016b) cast doubt on its validity and ascribe the "spanning puzzle" to model specification and measurement errors.

The current chapter adds to this line of term structure research by examining and quantifying the changes in the predictability in government bond markets and in the predictive content of macroeconomic variables when extending the sample period to the post-global financial crisis era.

The macroeconomic stability in the wake of the global financial crisis is also receiving increasing attention (Stock and Watson, 2012; Cheng et al., 2016; Aastveit et al., 2017; Clark and McCracken, 2017), while Giannone, Lenza, and Reichlin (2019) implement a Bayesian VAR (BVAR) model to explore the linkages between macroeconomic variables and bonds with specific maturities in the Eurozone. Complementing this strand of research, the current chapter provides insights into a similar matter in the U.S. economy with particular focus placed on the links between the term structure and macroeconomic fundamentals, by exploiting the capability of DFM in terms of characterising latent factors and dynamics.

Applying the conditional forecasting approach, which has demonstrated its usefulness in facilitating model validation (Stock and Watson, 2012; Cheng et al., 2016; Aastveit et al., 2017; Clark and McCracken, 2017; Debortoli et al., 2019), this chapter first constructs DFM-based conditional forecasts for the yields based on ex-post realised paths of macroeconomic indicators. Significant deviations of the model-implied paths from the actual realisations would signal either unstable macro-yields relationships or unusual yield shocks during and after the crisis. Compared to existing research, which builds conclusions primarily on graphical assessments of conditional forecasts, this chapter additionally conducts a formal out-of-sample forecast evaluation and examines the potential of this approach as a possible solution for yield forecasting during this prolonged period of low-interest rates.

3.3 Model Specification

3.3.1 Macro-Finance Term Structure Models

Researchers have proposed a variety of approaches to obtaining an explicit and consistent description of bonds across different maturities. The Nelson–Siegel approach (Nelson and Siegel, 1987), together with its dynamic version (Diebold and Li, 2006), can be interpreted as an interpolation and extrapolation tool that effectively captures the varying shapes of the yield curve. The Nelson–Siegel model enjoys widespread popularity due to its parsimonious specification and remarkable performance in fitting and forecasting. Consider a data set consisting of yield observations with different maturities $Y_t = (Y_{\tau_1,t}, \cdots, Y_{\tau_n,t})'$, the Nelson–Siegel approach validates the level, slope, and curvature interpretations of the latent factors $(X_{L,t}, X_{S,t}, X_{C,t})$ by explicitly specifying the following factor loading matrix

$$\begin{pmatrix} Y_{\tau_1,t} \\ Y_{\tau_2,t} \\ \cdots \\ Y_{\tau_n,t} \end{pmatrix} = \begin{pmatrix} 1 & \frac{1-e^{-\lambda\tau_1}}{\lambda\tau_1} & \frac{1-e^{-\lambda\tau_1}}{\lambda\tau_1} - e^{-\lambda\tau_1} \\ 1 & \frac{1-e^{-\lambda\tau_2}}{\lambda\tau_2} & \frac{1-e^{-\lambda\tau_2}}{\lambda\tau_2} - e^{-\lambda\tau_2} \\ & \cdots & \cdots \\ 1 & \frac{1-e^{-\lambda\tau_n}}{\lambda\tau_n} & \frac{1-e^{-\lambda\tau_n}}{\lambda\tau_n} - e^{-\lambda\tau_n} \end{pmatrix} \begin{pmatrix} X_{L,t} \\ X_{S,t} \\ X_{C,t} \end{pmatrix}. \tag{3.1}$$

The level factor $X_{L,t}$ influences the curve as a whole, corresponding to a constant loading for all maturities. The loading on the slope factor decays monotonically from unity to zero as the maturity τ extends, causing the yield curve to flatten when $X_{S,t}$ increases. The loading function on the curvature factor achieves its maximum value at intermediate durations, generating shifts with humpedness. To characterise the evolution of yields over time, the dynamic version of the Nelson–Siegel model (Diebold and Li, 2006) assumes that the underlying factors follow a univariate AR(1) or multivariate VAR(1) process

$$X_{t+1} = \mu + \Pi X_t + \epsilon_t. \tag{3.2}$$

Differing from the Nelson–Siegel approach in several important ways, the no-arbitrage term structure models are derived in a more theoretically rigorous fashion utilising the no-arbitrage argument. Assuming an exponentially affine pricing kernel in an arbitrage-free market, factor loadings are derived using recursive pricing formulas

$$\begin{pmatrix} Y_{\tau_1,t} \\ Y_{\tau_2,t} \\ \cdots \\ Y_{\tau_n,t} \end{pmatrix} = \begin{pmatrix} -\frac{1}{\tau_1}A(\tau_1) \\ \cdots \\ -\frac{1}{\tau_n}A(\tau_n) \end{pmatrix} + \begin{pmatrix} -\frac{1}{\tau_1}B(\tau_1)' \\ \cdots \\ -\frac{1}{\tau_n}B(\tau_n)' \end{pmatrix} \begin{pmatrix} X_{1,t} \\ \cdots \\ X_{m,t} \end{pmatrix}, \tag{3.3}$$

where $A(\tau)$ and $B(\tau)$ are functions of maturity τ and parameters that govern the risk-neutral dynamics of X_t. Despite the Nelson–Siegel approach itself not being necessarily associated with the notion of no-arbitrage, it could be specified in a manner that aligns with the affine term structure framework. This can be achieved by deploying certain dynamics for the underlying states X_t so that the arbitrage-free bond pricing formulas $A(\tau)$ and $B(\tau)$ in Equation (3.3) match the Nelson–Siegel representation in Equation (3.1), as suggested by the arbitrage-free Nelson–Siegel model (Christensen et al., 2009, 2011)

$$\begin{pmatrix} dX_{L,t} \\ dX_{S,t} \\ dX_{C,t} \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \lambda & -\lambda \\ 0 & 0 & \lambda \end{pmatrix} \left[\begin{pmatrix} \theta_L^Q \\ \theta_S^Q \\ \theta_C^Q \end{pmatrix} - \begin{pmatrix} X_{L,t} \\ X_{S,t} \\ X_{C,t} \end{pmatrix} \right] dt + \Sigma \begin{pmatrix} dW_{L,t}^Q \\ dW_{S,t}^Q \\ dW_{C,t}^Q \end{pmatrix}. \quad (3.4)$$

The drift vector θ and volatility matrix Σ in the factor dynamics (Equation (3.4)) result in a yield-adjustment term $C(\tau)/\tau$ that distinguishes the arbitrage-free Nelson–Siegel model from the dynamic Nelson–Siegel model

$$Y_{\tau,t} = X_{L,t} + \frac{1 - e^{-\lambda\tau}}{\lambda\tau} X_{S,t} + (\frac{1 - e^{-\lambda\tau}}{\lambda\tau} - e^{-\lambda\tau}) X_{C,t} - \frac{C(\tau)}{\tau}. \quad (3.5)$$

The macro-finance term structure models extend the conventional yields-based models by combining additional macroeconomic components into the factor set X. Consider a set of observations from the macroeconomic panel $Y_{M,t}$, the joint dynamics of the government bond yields and macroeconomic variables are addressed by augmenting the state vector with macro factors $X_{M,t}$

$$\begin{pmatrix} Y_{\tau,t} \\ Y_{M,t} \end{pmatrix} = \begin{pmatrix} c_Y \\ c_M \end{pmatrix} + \begin{pmatrix} \Gamma_{YY} & \Gamma_{YM} \\ \Gamma_{MY} & \Gamma_{MM} \end{pmatrix} \begin{pmatrix} X_{Y,t} \\ X_{M,t} \end{pmatrix} + \begin{pmatrix} e_{Y,t} \\ e_{M,t} \end{pmatrix}. \quad (3.6)$$

The dynamic factor approach further assumes that the common factors and idiosyncratic disturbances follow autoregressive processes

$$\begin{pmatrix} X_{Y,t} \\ X_{M,t} \end{pmatrix} = \begin{pmatrix} \mu_Y \\ \mu_M \end{pmatrix} + \begin{pmatrix} \Pi_{YY} & \Pi_{YM} \\ \Pi_{MY} & \Pi_{MM} \end{pmatrix} \begin{pmatrix} X_{Y,t-1} \\ X_{M,t-1} \end{pmatrix} + \eta_t, \eta_t \sim N(0, \Sigma_\eta) \quad (3.7)$$

$$e_t = \Pi_e e_{t-1} + v_t, v_t \sim N(0, \Sigma_v). \quad (3.8)$$

Rather than predetermining the number of latent states, Coroneo, Giannone, and Modugno (2016) select the optimal number of macro factors $X_{M,t}$ according to the modified Bai–Ng information criterion (Bai and Ng, 2002). A five-factor model, that is two additional macro factors combined with Nelson–Siegel yield factors identified by Equation (3.1), captures the dynamics of their sample data (1970:M1–2008:M12).

3.3.2 Unspanning Restrictions

An important question raised in earlier studies concerns the wedge between factors spanning the cross section of contemporaneous yields and those driving future yields (Ludvigson and Ng, 2009; Duffee, 2011b; Wright, 2011; Bauer et al., 2014; Joslin et al., 2014; Cieslak and Povala, 2015; Bauer and Rudebusch, 2016b; Feldhütter et al., 2016). In order to exploit the yield information contained in macroeconomic activities, the standard macro-finance term structure models extend the traditional term structure models by directly incorporating macro factors into the latent factor space that determines short rates and pricing kernels. The invertibility of these models implies that the macro factors can be expressed as linear combinations of the current yields, that is, the macroeconomic panel can be fully spanned by the contemporaneous yield curve. Some research, however, finds evidence against this spanning property, as indicated by the low R-squared values produced by regressions of macroeconomic variables on yields. Accordingly, researchers are proposing a new class of unspanned macro-finance term structure models that do not impose this invertibility, while still incorporating the predictive content embedded in the macroeconomic panel beyond that already provided by the cross section of yields (Cooper and Priestley, 2008; Ludvigson and Ng, 2009; Greenwood and Vayanos, 2014).

This new class of macro-finance term structure models specifies the macro factors to be unspanned in such a manner whereby they exert no contemporaneous effects on the yield curve, but interact with the yield factors to drive the evolution of the yield curve. Within the framework adopted by Coroneo, Giannone, and Modugno (2016), this unspanning feature can be easily achieved by imposing restrictions on the observation matrix Γ and the state transition matrix Π. With the factor loadings of yields on macro factors Γ_{YM} constrained to zero, the macro factors $X_{M,t}$ bear no relation to the contemporaneous yields $Y_{\tau,t}$, but affect the future yields $Y_{\tau,t+h}$ through their interactions with the future yield factors $X_{Y,t+h}$.

There are two aspects to the spanning puzzle: i) whether the macro components lie in the span of the contemporaneous yield curve, captured by the factor loadings of current yields on current macro factors Γ_{YM} ; and, ii) whether they have predictive content for future yields, characterised by the factor loadings of the future yield factors on the current macro factors Π_{YM}. Both properties can be checked individually by examining the following restrictions through likelihood ratio tests:

$$\text{i)} \quad H_0 : \Gamma_{YM} = 0, \quad H_1 : \Gamma_{YM} \neq 0; \tag{3.9}$$

$$\text{ii)} \quad H_0 : \Pi_{YM} = 0, \quad H_1 : \Pi_{YM} \neq 0. \tag{3.10}$$

The acceptance of the first null hypothesis that macro factors are not fully spanned by the contemporaneous yield curve ($\Gamma_{YM} = 0$), along with the rejection of the second null hypothesis, which confirms the predictive value of the macro elements ($\Pi_{YM} \neq 0$), provides solid evidence for the unspanning restrictions. Combining the information criteria and likelihood ratio tests, Coroneo, Giannone, and Modugno (2016) choose the five-factor model with unspanning restrictions over alternative model specifications.

3.3.3 Expectation Maximisation Estimation and Conditional Forecasting

Augmenting the underlying states with the idiosyncratic components casts the DFM specified by Equations (3.6–3.8) into the standard state-space representation

$$Y_t = \Gamma^* X_t^* + e_t^*, e_t^* \sim N(0, R^*), \tag{3.11}$$

$$X_t^* = \Pi^* X_{t-1}^* + \eta_t^*, \eta_t^* \sim N(0, Q^*), \tag{3.12}$$

which is amenable to efficient implementation using the expectation maximisation algorithm coupled with filtering and smoothing techniques. For simplicity, the asterisks in Equation (4.10) and Equation (4.11) are omitted in this section. Given the parametric specification, the log-likelihood of the observed data $\{y_t\}_{t=1\ldots n}$ and underlying states $\{x_t\}_{t=1\ldots n}$ is

$$\begin{aligned}\log L = &-\frac{1}{2} n \log |R| - \frac{1}{2} \sum_{t=1}^{n} (y_t - \Gamma x_t)^T R^{-1} (y_t - \Gamma x_t) \\ &-\frac{1}{2} n \log |Q| - \frac{1}{2} \sum_{t=1}^{n} (x_t - \Pi x_{t-1})^T Q^{-1} (x_t - \Pi x_{t-1}) \\ &-\frac{1}{2} \log |\Sigma_{x_0}| - \frac{1}{2} (x_0 - E(x_0))^T \Sigma_{x_0}^{-1} (x_0 - E(x_0)),\end{aligned} \tag{3.13}$$

with x_0 and Σ_{x_0} referring to the initial state and covariance. Let $x_{i|j}$ be the conditional expectation of state x at time i based on the information available at time j, that is $\hat{x}_{i|j} = E(x_i|\Omega_j)$, and $P_{i|j}$ be the error covariance matrices $P_{i|j} = \mathrm{Var}(x_i|\Omega_j)$ and $P_{i,i-1|j} = \mathrm{Cov}(x_i, x_{i-1}|\Omega_j)$, maximising the conditional expectation of the above log-likelihood given full observations Ω_T is equivalent to solving the following optimisation problem

$$\max E(\log L|\Omega_T) = -\frac{1}{2}n\log|Q| - \frac{1}{2}tr\{Q^{-1}(C - B\Pi^T - \Pi B^T + \Pi A\Pi^T)\} - \frac{1}{2}n\log|R|$$
$$- \frac{1}{2}tr\{R^{-1}(\sum_{t=1}^{n}(y_t - \Gamma x_{t|T})(y_t - \Gamma x_{t|T})^T + \sum_{t=1}^{n}\Gamma P_{t|T}\Gamma^T)\}$$
$$- \frac{1}{2}\log|\Sigma_{x_0}| - \frac{1}{2}tr\{\Sigma_{x_0}^{-1}(P_{0|T} + (x_{0|T} - E(x_0))(x_{0|T} - E(x_0))^T)\}, \tag{3.14}$$

where $A = \sum_{t=1}^{n}(x_{t-1|T}x_{t-1|T}^T + P_{t-1|T})$, $B = \sum_{t=1}^{n}(x_{t|T}x_{t-1|T}^T + P_{t,t-1|T})$, and $C = \sum_{t=1}^{n}(x_{t|T}x_{t|T}^T + P_{t|T})$. The conditional mean $x_{t|T}$ and corresponding covariance matrices $P_{t|T}$ and $P_{t,t-1|T}$ are generated by running the Kalman filter and smoother through the state-space structure.

Progressing forward in time, the Kalman filter updates estimates as new observations arrive by first generating projections of the last posteriors $\hat{x}_{t-1|t-1}$ through the state transition matrix Π

$$\hat{x}_{t|t-1} = \Pi\hat{x}_{t-1|t-1}, \tag{3.15}$$
$$P_{t|t-1} = \Pi P_{t-1|t-1}\Pi^T + Q, \tag{3.16}$$

and then blending these prior estimates $\hat{x}_{t|t-1}$ with direct measurements y_t by means of the Kalman correction term K_t, which yields the posterior state $\hat{x}_{t|t}$ and covariance $P_{t|t}$

$$K_t = P_{t|t-1}\Gamma^T(\Gamma P_{t|t-1}\Gamma^T + R)^{-1}, \tag{3.17}$$
$$\hat{x}_{t|t} = \hat{x}_{t|t-1} + K_t(y_t - \Gamma\hat{x}_{t|t-1}), \tag{3.18}$$
$$P_{t|t} = P_{t|t-1} - K_t\Gamma P_{t|t-1}. \tag{3.19}$$

In the backward pass, the Kalman smoother starts from the final timestep $\hat{x}_{T|T}$ and adjusts all posterior estimates $\hat{x}_{t|t}$ and $P_{t|t}$ by the deviations of the smoothed estimates $\hat{x}_{t+1|T}$ from the prior estimates $\hat{x}_{t+1|t}$ (Ansley and Kohn, 1982)

$$\hat{x}_{t|T} = \hat{x}_{t|t} + P_{t|t}\Pi^T P_{t+1|t}^{-1}(\hat{x}_{t+1|T} - \hat{x}_{t+1|t}), \tag{3.20}$$
$$P_{t|T} = P_{t|t} + P_{t|t}\Pi^T P_{t+1|t}^{-1}(P_{t+1|T} - P_{t+1|t})(P_{t|t}\Pi^T P_{t+1|t}^{-1})^T. \tag{3.21}$$

Based on the conditional distributions produced by the above Kalman procedures, which constitute the expectation step of the expectation maximisation algorithm, the maximisation step obtains the parameter values by solving Equation (3.14)

subject to linear constraints $L_m \text{ vec } \Gamma = m$ and $L_s \text{ vec } \Pi = s$,

$$\text{vec } \Pi = \text{vec}(BA^{-1}) + (A^{-1} \otimes Q)L_s^T\{L_s(A^{-1} \otimes Q)L_s^T\}^{-1}\{s - L_s\text{vec}(BA^{-1})\}, \tag{3.22}$$

$$Q = n^{-1}(C - B\Pi^T - \Pi B^T + \Pi A \Pi^T), \tag{3.23}$$

and similarly for Γ and R (Wu et al., 1996). Both steps are recursively implemented until convergence, rendering consistent estimates for the states and parameters.

The Kalman filter smoother, furthermore, enables the plausible treatment of missing values and the generation of conditional forecasts. For ease of illustration, let y_t denote the information available at time t, that is $y_t = \{y_1, \ldots, y_t\}$. The Kalman smoother refines the posteriors generated in the forward pass $\hat{x}_{t|t}$ using the full set of observations

$$\hat{x}_{t|T} = E[x_t|\Omega_T] = E[x_t|y_T, \Theta]. \tag{3.24}$$

If a variable $y^{i\in\mathcal{I}}$ is unobservable over the interval $[t_C, T]$, the estimates $\hat{y}_{t|T}^{i\in\mathcal{I}}$ generated by the Kalman smoother are essentially the projections of its own observations up to t_C and the entire paths of the other variables

$$\hat{y}_{t|T}^{i\in\mathcal{I}} = E[y_t^{i\in\mathcal{I}}|\Omega_T] = E[y_t^{i\in\mathcal{I}}|y_{t_C}^{i\in\mathcal{I}}, y_T^{i\notin\mathcal{I}}]. \tag{3.25}$$

As suggested by Camba-Mendez (2012), this can be achieved by adjusting the optimal gain K_t over the period $[t_C, T]$ using

$$K_t = P_{t|t-1}\Gamma^T(G\Gamma P_{t|t-1}\Gamma^T + GR)^{-1}, \tag{3.26}$$

with G indicating where observations are available. This is equivalent to the modified state-space representation approach in Bańbura, Giannone, and Lenza (2015) that directly removes the rows and columns in Γ and R that correspond to the missing data at time t. Implementing the standard Kalman procedures generates posterior estimates

$$\hat{y}_{t|t} = E[y_t|y_t^{i\in\mathcal{I}}, y_t^{i\notin\mathcal{I}}], \text{when } t < t_C, \tag{3.27}$$

$$\hat{y}_{t|t} = E[y_t|y_{t_C}^{i\in\mathcal{I}}, y_t^{i\notin\mathcal{I}}], \text{when } t > t_C, \tag{3.28}$$

with the Kalman smoother then refining these estimates backwards from T to incorporate all the available information

$$\hat{y}_{t|T} = E[y_t|\Omega_T] = E[y_t|y_{t_C}^{i\in\mathcal{I}}, y_T^{i\notin\mathcal{I}}]. \tag{3.29}$$

When observations are missing for variable $i \in \mathcal{I}$, Equation (3.29) provides projections based on the entire paths of the other variables $y_T^{i\notin\mathcal{I}}$ and its own observations up to t_C, that is $y_{t_C}^{i\in\mathcal{I}}$.

Thus far, the Kalman method lays out an insightful framework for handling missing data, which in the context of term structure modelling can be exploited to construct counterfactual nominal rates in the low interest rate environment. It also provides a valuable avenue for examining the dynamic interactions among macroeconomic variables and the Treasury yield curve, after controlling for the effects of anomalous behaviour of specific yields. Section 3.5.1 conducts a counterfactual scenario assessment by treating a subset of yield variables as unobservable post-2008:M10 (following Bańbura, Giannone, and Lenza (2015)), and inserting the missing values via the conditional forecasts implied by Equation (3.29), where the full information set includes observations of these yields Y^{τ_i} up to time t_C =2008:M10 and the full paths of other variables up to time T =2016:M12, that is

$$\Omega = \{Y_{2008:\text{M10}}^{\tau_i}, Y_{2016:\text{M12}}^{\tau_j\neq\tau_i,M}\}. \tag{3.30}$$

3.4 Predictability and Underlying Causes

This section examines the performance of the unspanned macro-yields model proposed by Coroneo, Giannone, and Modugno (2016) and other relevant approaches in an updated sample spanning 1970:M1–2016:M12 to cover the crisis and subsequent recovery, since when the low interest rate has become a much more significant issue. This section then identifies the underlying causes of performance deterioration by decomposing the realised forecast errors into those of the latent factors, and by conducting generalised impulse response analysis in order to examine the propagation of shocks among the cross section of yields and the macroeconomic panel.

Over the extended timespan 1970:M1–2016M:12, the results of the in-sample analysis moderately change, retaining similar findings with weaker statistical significance or achieving slightly different conclusions. Despite the information criteria still favouring the inclusion of macroeconomic variables, the five-factor model no longer exhibits noticeable superiority over other multi-factor models, and the in-sample predictability of the macro-related models tends to deteriorate with the updated sample. With respect to the out-of-sample analysis, forecasts are generated recur-

sively using data from 1970:M1 and evaluated over the 1990:M1–2016:M12 period. The macro-yields model suffers from a more pronounced deterioration in terms of the out-of-sample forecasting performance, with the additional sample constituting a sizeable fraction of the evaluation period, that is 96 out of 324 observations.

3.4.1 In-Sample and Out-of-Sample Predictability

In accordance with the paper by Coroneo, Giannone, and Modugno (2016), the yield dataset utilised is the Fama–Bliss dataset from the Centre for Research in Security Prices, covering 3-month, 1-, 2-, 3-, 4-, 5-year zero-coupon bond yields. Macroeconomic data are primarily collected from the Federal Reserve Economic Data (FRED), consisting of 14 leading economic indicators that reflect various aspects of the macroeconomy. The Appendices provide detailed variable descriptions and sources, and an implementation in the original 1970:M1–2008:M12 timespan.

As seen in Figure 3.1, the yields exhibit small-scale fluctuations in the post-crisis period relative to the pre-2009 data, particularly for the short-term rates. By comparing the yield dynamics with those of the macroeconomic variables in Figure 3.2, we can easily observe that bond yields and macroeconomic variables behave distinctly over the extended period. In the post-global financial crisis, denoted by the grey-shaded area, the industrial production (IP) growth and producer price index (PPI) growth revert to their historical ranges in a relatively short period of time, fluctuate moderately and then decline gradually, while the yields are stuck at unprecedented low values. The exceedingly low level of variation in these independent variables undermines the reliability of the macro-yields model. Notwithstanding the strong tendency of macroeconomic indicators to return to their previous paths in the post-2009 recovery, the yield curve continues its sluggish movement, posing a severe challenge to macro-finance modelling during this period.

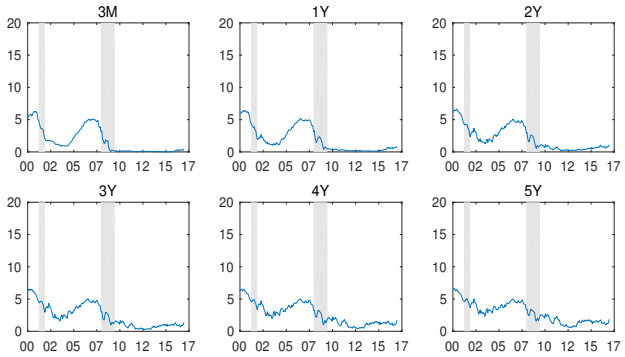

Figure 3.1: Time series of the yields. The figure plots the 3m-, 1y-, 2y-, 3y-, 4y-, and 5y Fama–Bliss zero-coupon yields over the period 2000:M1–2016:M12. The grey-shaded areas mark the contraction periods identified by NBER.

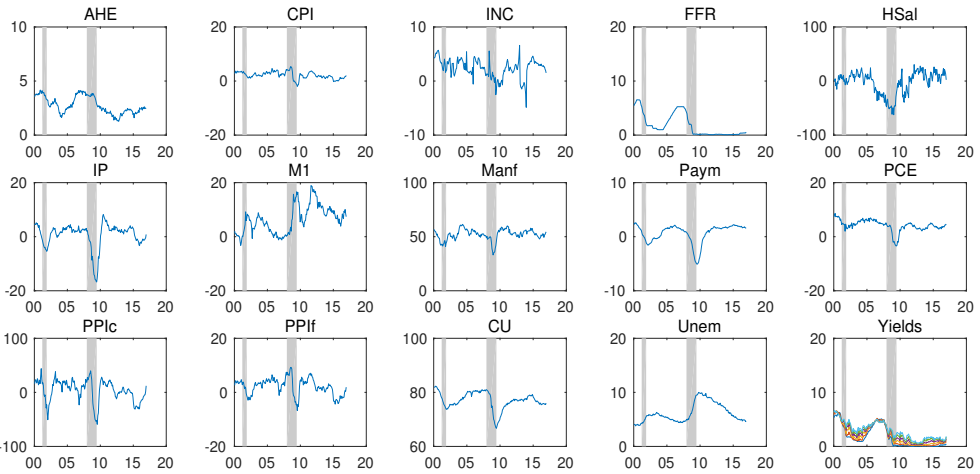

Figure 3.2: Time Series of Macroeconomic Variables. The figure plots the time series of macroeconomic variables over the period 2000:M1–2016:M12. Please see Appendix 3.A for further details on the data and variables.

Evidence favouring selection of the five-factor model in Table 3.1 is currently weaker, with the information criterion being indistinguishable from that of the four-factor model. In contrast to the original sample over which the five-factor model yields the lowest information criterion of -0.11, the extended sample no longer represents conclusive evidence supporting the unspanned five-factor model in the post-2008 data, as suggested by the negligible differences in terms of the information criterion values.[2]

Table 3.1: Selecting the number of macro factors

	1970:M1 – 2008:M12		1970:M1 – 2016:M12	
Number of factors	IC	V	IC	V
3	0.02	0.44	0.05	0.45
4	**-0.03**	0.31	**0.00**	0.32
5	**-0.11**	0.22	**0.00**	0.24
6	0.01	0.18	0.23	0.23
7	0.23	0.17	0.39	0.20
8	0.42	0.16	0.37	0.15

Note. This table reports the modified information criterion (IC) that accounts for both the cross-sectional dimension and the time-series dimension (Bai and Ng, 2002), which is calculated from the average residual variance V. Results are reported for k-factor (k=3,4,5,6,7,8) models in rows, with the left panel corresponding to the original sample (1970:M1–2008:M12) and the right panel to the extended sample (1970:M1–2016:M12).

To a large extent, the yield curve factors still explain the variation in yields, at over 99% for government bond yields and 97% for the effective federal funds rate

[2]Caruso and Coroneo (2019) have recently reached a similar conclusion utilising a mixed-frequency model and a different dataset.

(FFR), as shown in Table 3.2. Similar to the results reported by Coroneo, Giannone, and Modugno (2016), the unspanned macro factors are closely associated with the IP and inflation-related indicators (e.g. CPI, PPI). With the incorporation of the first unspanned macro factor UM1, the cumulative variance explained by the macro-yield factors increases by 69% for IP, 60% for total nonfarm payrolls (Paym), and 53% for the PMI composite (Manf) index, indicating substantial relationships between the first unspanned macro factor and the real economic activity. The introduction of the second unspanned macro factor UM2 increases the proportion of variance explained in the producer price index on finished goods (PPIf) by 48% and in the consumer price index (CPI) by 27%, relating the second unspanned factor to measures of inflation.

Table 3.2: Cumulative variance explained by the macro-yields factors

	1970:M1 – 2008:M12					1970:M1 – 2016:M12				
	Level	Slope	Curv	UM1	UM2	Level	Slope	Curv	UM1	UM2
3m	0.59	0.94	1.00	1.00	1.00	0.60	0.93	1.00	1.00	1.00
1y	0.61	0.83	0.99	0.99	0.99	0.59	0.81	0.99	0.99	0.99
2y	0.65	0.78	0.99	0.99	0.99	0.62	0.76	1.00	1.00	1.00
3y	0.70	0.79	1.00	1.00	1.00	0.65	0.76	1.00	1.00	1.00
4y	0.74	0.80	0.99	0.99	0.99	0.69	0.77	1.00	1.00	1.00
5y	0.78	0.82	0.99	0.99	0.99	0.73	0.80	1.00	1.00	1.00
AHE	0.07	0.29	0.33	0.33	0.67	0.14	0.37	0.43	0.43	0.60
CPI	0.19	0.49	0.48	**0.50**	**0.85**	0.26	0.59	0.56	**0.56**	**0.83**
INC	0.00	0.02	0.03	0.34	0.36	0.00	0.01	0.06	0.33	0.41
FFR	0.53	0.93	0.96	0.96	0.97	0.56	0.94	0.97	0.97	0.97
Hsal	0.00	0.20	0.19	0.23	0.23	0.00	0.14	0.12	0.18	0.19
IP	0.02	0.02	**0.03**	**0.70**	0.70	0.00	0.00	**0.02**	**0.71**	0.71
M1	0.17	0.25	0.26	0.25	0.31	0.01	0.07	0.11	0.14	0.16
Manf	0.03	0.06	**0.05**	**0.61**	0.65	0.02	0.06	**0.04**	**0.57**	0.59
Paym	0.00	0.02	**0.10**	**0.71**	0.70	0.00	0.02	**0.11**	**0.71**	0.70
PCE	0.16	0.23	0.33	0.46	0.78	0.24	0.32	0.48	0.64	0.77
PPIc	0.03	0.14	0.14	0.20	0.43	0.00	0.10	0.08	0.29	0.57
PPIf	0.03	0.34	0.32	**0.33**	**0.80**	0.07	0.38	0.32	**0.32**	**0.80**
CU	0.02	0.16	0.21	0.63	0.64	0.01	0.14	0.25	0.70	0.70
Unem	0.44	0.54	0.55	0.65	0.68	0.12	0.18	0.18	0.39	0.43

Note. In each row, this table reports the cumulative variance of each yield or macro variable explained by the yield curve factors (Level, Slope, and Curv) and the unspanned macro factors (UM1 and UM2) over the original sample (left panel) and over the extended sample (right panel). The grey-shaded areas highlight rows that exhibit notable changes over the extended sample.

Notable changes occur in a few macroeconomic variables over the extended sample. The cumulative variance of the producer price index on crude materials (PPIc) significantly improves, while those of M1 money stock (M1) and unemployment rate (Unem) considerably deteriorate. Strongly associated with vital macroeconomic indicators, the underlying macro factors experience dramatic changes that

correspond with the economic downturn during the 2007–09 financial crisis. For the macroeconomic variables whose dynamic features are well captured by the macro factors in the pre-2009 data and parallel shift downward over the extended timespan, such as the PPIc, the explanatory power strengthens in the full sample. Meanwhile, for those that are not closely associated with the macro factors in the pre-2009 data and move in the opposite direction from the general trend, the explanatory power shrinks, as observed in M1 and house sales (HSal). For Unem, the only macroeconomic variable whose variation is explained to a greater extent by the yield factors (55%) than the macro factors (13%), its increase during the recession does not correspond with the sluggish movement of yields, which is constrained from below, leading to weaker explanatory power.

Table 3.3 summarises the in-sample predictability of the 12-month-ahead 1-year-holding-period excess bond returns over the entire 1970:M1–2016:M12 sample period. With the updated sample, nearly all R^2s decline considerably, with the most notable differences occurring in the macro-yields model where the R^2 decreases by 15% across maturities, pointing to a substantial drop in model performance. By comparing across models, the evidence supporting the outperformance of the macro-yields model over other models does not carry over into the 4- and 5-year maturities for which the LN+CP approach provides the best R^2. During the global financial crisis and its aftermath, the remarkable fluctuations that the macroeconomic indicators experience are disproportionately large compared to the yield movement, the magnitude of which is bounded from above. The deterioration is most pronounced in the macro-yields model that establishes an explicit relationship between the macro and yield movements employing the VAR methodology.

Table 3.3: In-Sample predictability of excess bond returns

	1970:M1 – 2008:M12					1970:M1 – 2016:M12				
Maturity	MY	OY	CP	LN	LN+CP	MY	OY	CP	LN	LN+CP
2y	0.55	0.12	0.22	0.33	0.41	0.40	0.12	0.16	0.22	0.30
3y	0.53	0.12	0.24	0.33	0.42	0.38	0.13	0.18	0.25	0.34
4y	0.49	0.14	0.27	0.32	0.43	0.35	0.15	0.21	0.25	0.36
5y	0.45	0.15	0.24	0.30	0.40	0.31	0.16	0.19	0.25	0.34

Note. This table reports the proportion of the variation in 1-year-ahead 1-year-holding-period excess returns accounted for by the macro-yields model (MY), the only-yields model (OY), the CP factor (variables constructed using 12-, 24-, 36-, 48-, 60-month forward rates), the LN factor (combinations of principal components extracted from 131 macroeconomic data series), and the joint LN+CP factor. The rows correspond to excess returns on 2- to 5-year bonds. The left panel shows the pre-2009 results, and the right panel displays the results of the extended sample.

Table 3.4 conducts tests of the unspanning conditions from two perspectives: whether yields with different maturities load on the macro factors ($H_0 : \Gamma_{yx} = 0$

in Equation (3.9)), and whether the macro factors convey additional predictive information regarding yield dynamics ($H_0 : \Pi_{yx} = 0$ in Equation (3.10)). In the extended sample, the rejection of the latter null hypothesis lends strong support to the predictive scope of macro factors ($\Pi_{yx} \neq 0$), while the test regarding the former hypothesis generates negative values.

Table 3.4: Likelihood ratio test for the unspanning restrictions

	1970:M1 – 2008:M12		1970:M1 – 2016:M12	
H_0	Test statistic	p-value	Test statistic	p-value
Γ_{yx}=0	12.85	0.38	-271.42	-
Π_{yx}=0	79.03	0.00	114.59	0.00

Note. The table reports the likelihood ratio test statistics and p-values for the null hypotheses $\Gamma_{yx} = 0$ in Equation (3.9) and $\Pi_{yx} = 0$ in Equation (3.10), with the left panel referring to the original sample and the right panel to the extended sample.

To further explain the findings of the model selection and likelihood ratio test, Figure 3.3 performs the PCA, utilising a 10-year rolling window, on the residuals from the regression of macroeconomic variables on the Nelson–Siegel yield factors. The residuals represent the unspanned component of the macroeconomic information, and the resultant principal components serve as reasonable proxies for the unspanned macro factors.

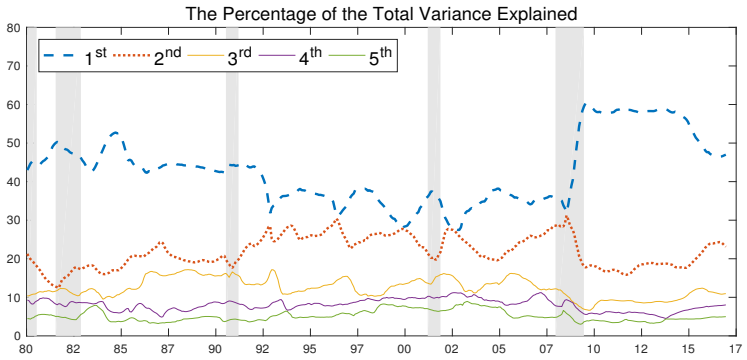

Figure 3.3: Rolling principal component analysis of unspanned macroeconomic information. The principal component analysis is conducted on the unspanned macroeconomic information utilising a 10-year rolling window, the end dates of which are reported along the x-axis. The figure reports the percentage of the total variance explained by the first five principal components. The grey-shaded areas indicate the contraction periods determined by the NBER.

Over the period of relative calm from the early 1990s to the mid-2000s, the first principal component (dashed blue line in Figure 3.3) explains approximately 30%–40% of the observed variation in the unspanned component, while the second principal component (dotted red line) accounts for around 20%–30% of the variation. During the 2007–09 period, with the onset of the financial crisis, the percentage

explained by the first principal component dramatically increases to approximately 60%, which is virtually the sum of those captured by the primary two principal components during the late 1990s and early 2000s. That is to say, since the recession, the first principal component accounts for the amount of variation that was used to be captured by two unspanned factors before the recession.

While the rMSFEs of the only-yields model slightly deteriorate in the extended sample, the performance of the macro-yields model relative to the random walk substantially worsens with all rMSFEs increasing to values above one (upper-right panel in Table 3.5), indicating generally inferior predictability for all maturities at all horizons. Long-horizon forecasting experiences the most significant deterioration, with rMSFEs increasing by an average value of 0.65 for forecasts at 12 and 24 months ahead and 0.37 across all horizons. Of particular note is that, in the extended sample, the only-yields model even delivers better forecasting performance than the macro-yields model on average, greatly invalidating the argument that favours incorporation of macroeconomic information into the Nelson–Siegel framework. In agreement with prevailing beliefs, the random walk forecasting remains a tough benchmark to beat.

Table 3.5: Out-of-sample predictability of the yield curve

	1970:M1 – 2008:M12						1970:M1 – 2016:M12					
					Macro-Yields							
Maturity	3m	1y	2y	3y	4y	5y	3m	1y	2y	3y	4y	5y
h=1	0.82	0.92	1.00	1.00	0.99	0.99	1.08	1.11	1.12	1.10	1.07	1.06
h=3	0.79	0.92	0.98	0.96	0.98	0.98	1.23	1.24	1.23	1.19	1.17	1.16
h=6	0.79	0.89	0.94	0.93	0.93	0.94	1.35	1.35	1.37	1.34	1.31	1.29
h=12	0.69	0.75	0.79	0.80	0.81	0.81	1.34	1.33	1.39	1.43	1.44	1.41
h=24	0.61	0.65	0.72	0.78	0.83	0.90	1.19	1.19	1.35	1.51	1.62	1.74
					Only-Yields							
Maturity	3m	1y	2y	3y	4y	5y	3m	1y	2y	3y	4y	5y
h=1	0.93	1.06	1.16	1.10	1.07	1.07	0.94	1.06	1.15	1.11	1.06	1.05
h=3	0.96	1.14	1.21	1.14	1.11	1.11	0.98	1.12	1.21	1.16	1.10	1.09
h=6	0.99	1.18	1.26	1.22	1.17	1.17	1.04	1.18	1.28	1.24	1.16	1.14
h=12	1.05	1.16	1.26	1.27	1.25	1.25	1.09	1.19	1.31	1.33	1.28	1.25
h=24	1.09	1.14	1.28	1.40	1.49	1.61	1.17	1.21	1.40	1.52	1.57	1.63

Note. The table reports the rMSFE of the macro-yields model (upper panel) and the only-yields model (lower panel) with respect to the random walk for forecasts of 3m-, 1y-, 2y-, 3y-, 4y-, and 5y-yields (in columns) made h=1, 3, 6, 12, 24 months ahead (in rows). The left panel shows the pre-2009 results, and the right panel displays the results of the extended sample.

A similar pattern is observed in Table 3.6, which examines the 1-year-holding-period excess bond returns. Extending the sample into the post-global financial crisis era increases the rMSFEs with respect to the expectations hypothesis by 0.55

on average across different maturities. Table 3.6 further suggests that both the macro-yields model and the only-yields model cannot surpass the expectations hypothesis over the extended period, as indicated by the rMSFEs with values higher than one. Amongst the models considered, the LN+CP approach delivers the most accurate forecasts and represents the only model that can compete with the expectations hypothesis benchmark. While introducing macro factors into the only-yields model produces inferior results over the 1970:M1–2016:M12 period, combining the macroeconomic variables (LN) into the set of yield factors (CP) improves the forecast accuracy. This observation suggests that the macroeconomic panel does convey additional insights into the future development of yields, but the way in which the macroeconomic information is exploited requires careful consideration, especially during the recession and recovery periods.

Table 3.6: Out-of-sample predictability of excess bond returns

	1970:M1 – 2008:M12					1970:M1 – 2016:M12				
Maturity	MY	OY	CP	LN	LN+CP	MY	OY	CP	LN	LN+CP
2y	0.77	1.20	1.17	0.79	0.80	1.37	1.23	1.32	1.22	0.96
3y	0.75	1.19	1.20	0.78	0.83	1.30	1.23	1.40	1.07	0.90
4y	0.75	1.18	1.20	0.78	0.83	1.27	1.18	1.41	0.97	0.84
5y	0.76	1.17	1.18	0.80	0.83	1.28	1.13	1.40	0.90	0.82

Note. The table reports the rMSFE of different approaches with respect to the expectations hypothesis for forecasts of 1-year-holding-period excess returns on 2- to 5-year bonds (in rows). The columns correspond to the macro-yields model (MY), the only-yields model (OY), the CP factor, the LN factor, and the joint LN+CP factor, respectively. The left panel shows the pre-2009 results, and the right panel displays the results of the extended sample.

Tracking model performance over time, Figure 3.4 and Figure 3.5 revisit the squared forecast errors, as displayed in Coroneo, Giannone, and Modugno (2016), of bond yields and 1-year-holding-period excess bond returns at the 12-month-ahead forecast horizon, averaged over a 5-year rolling window. The performance of the macro-yields model does not perform well in the post-recession period. The smoothed forecast errors of the macro-yields model (the continuous blue line in Figure 3.4) trend upwards in the wake of the recession, spiking to cross the random walk curve (the dashed red line) from below in early 2010 and far outstripping it thereafter. Despite the crossover that indicates indifference between the macro-yields model and the random walk approach occurring around 2010, the macro-yields model begins to deliver inferior forecasts before the crossover point after taking into account the smoothing effect. The underperformance of the only-yields model relative to the random walk carries over into the extended evaluation period.

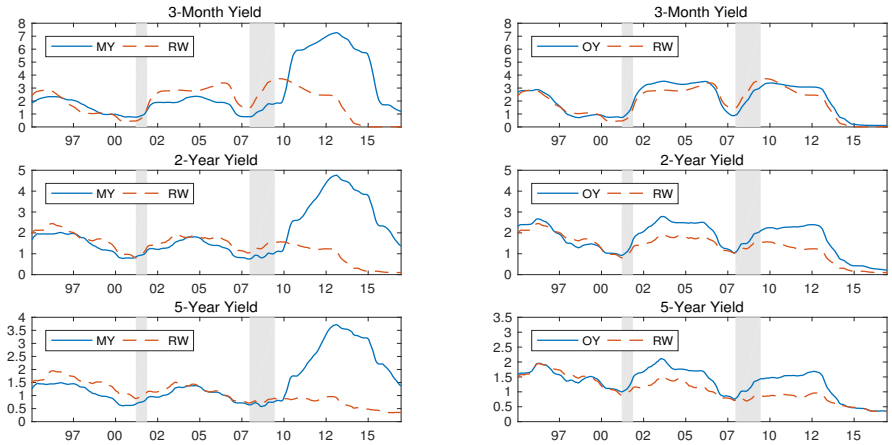

Figure 3.4: Forecasting the yield curve over time (smoothed squared errors). This figure displays the smoothed squared errors of 1-year-ahead forecasts for 3-month (upper), 2-year (middle), and 5-year (lower) yields over a 5-year rolling window, with the end dates of rolling windows reported along the x-axis. The continuous lines plot the macro-yields model (MY) and the only-yields model (OY), with the dashed lines representing the random walk benchmark (RW). The grey-shaded areas indicate the contraction periods determined by the NBER.

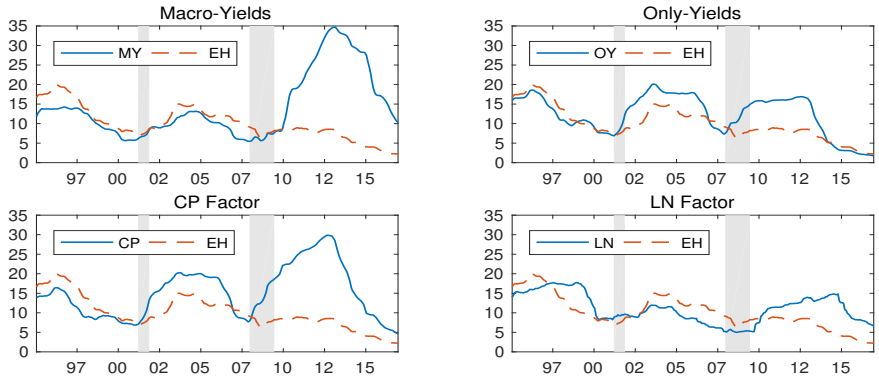

Figure 3.5: Forecasting excess bond returns over time (smoothed squared errors). This figure displays the smoothed squared errors of 1-year-ahead forecasts for 1-year-holding-period excess returns over a 5-year rolling window, with the end dates of rolling windows reported along the x-axis. The continuous lines plot the macro-yields model (MY), the only-yields model (OY), the LN factor, the CP factor, respectively, with the dashed lines reporting the expectations hypothesis benchmark (EH). The grey-shaded areas indicate the contraction periods determined by the NBER.

A similar pattern can be observed for the smoothed forecast errors for the 12-month-ahead 1-year-holding-period excess bond returns (see Figure 3.5). Of particular note is the outperformance of the expectations hypothesis over all models in the post-recession period. To conclude, the extended sample does not resoundingly support, and sometimes even invalidates the five-factor macro-yields model with unspanning restrictions.

3.4.2 Real-Time Performance

Macroeconomic data is usually subject to multiple rounds of revisions, which might affect and undermine the reliability of conclusions based upon them (Croushore and Stark, 2001; Garratt and Vahey, 2006).[3] To examine the extent to which the data revisions over time affect the conclusions on predictability, this subsection repeats the out-of-sample analysis using the real time data in Table 3.7 and Table 3.8.

Table 3.7: Out-of-sample forecast for the yield curve in real time

	Vintage Data						Real-Time Data					
	1970:M1 – 2008:M12											
Maturity	3m	1y	2y	3y	4y	5y	3m	1y	2y	3y	4y	5y
h=3	0.79	0.92	0.98	0.96	0.98	0.98	0.83	0.95	1.01	0.99	1.00	1.00
h=6	0.79	0.89	0.94	0.93	0.93	0.94	0.82	0.93	0.98	0.96	0.97	0.97
h=12	0.69	0.75	0.79	0.80	0.81	0.81	0.71	0.77	0.82	0.83	0.84	0.83
h=24	0.61	0.65	0.72	0.78	0.83	0.90	0.63	0.66	0.74	0.80	0.84	0.91
Maturity	1970:M1 – 2016:M12											
h=3	1.23	1.24	1.23	1.19	1.17	1.16	1.19	1.21	1.22	1.18	1.16	1.15
h=6	1.35	1.35	1.37	1.34	1.31	1.29	1.27	1.30	1.35	1.33	1.30	1.27
h=12	1.34	1.33	1.39	1.43	1.44	1.41	1.24	1.25	1.34	1.39	1.40	1.36
h=24	1.19	1.19	1.35	1.51	1.62	1.74	1.10	1.13	1.30	1.46	1.56	1.66

Note. The table reports the rMSFE of the macro-yields model (MY) with respect to the random walk for forecasts of 3m-, 1y-, 2y-, 3y-, 4y-, and 5y-yields (in columns) made h=3, 6, 12, 24 months ahead (in rows). The upper panel corresponds to the original timespan 1970:M1–2008:M12 while the lower panel corresponds to the extended timespan 1970:M1–2016:12. The table shows the results using vintage (left panel) and real-time data (right panel), respectively.

For the original timespan 1970:M1–2008:M12, using real-time data increases the relative MSFE by 0.02 on average (upper panel in Table 3.7), indicating slightly weaker outperformance of macro-yields model over the random walk. The deterioration in forecast performance with real-time data weakens the previous argument marginally, with the improvement at shorter horizons not as distinguishable as with the vintage data. The performance of macro-yields model during the post-global financial crisis period, however, gets slightly better with real-time data than with vintage data, especially for bonds with shorter maturities (lower panel).

Overall, the out-of-sample predictability of the yield curve and excess bond returns over the extended period substantially deteriorates (also see Table 3.8), whether with the vintage data or real-time data, leaving the qualitative nature of

[3]Retrieving macroeconomic data as the vintage versions of economic dataset available from St. Louis Fed's Economic Research Division, the original dataset used by Coroneo, Giannone, and Modugno (2016) matches well the historical vintages 2012:M5 and 2012:M6. Accordingly, this chapter uses the latest available data when the work was first conducted, the 2017:M5 vintage data.

the results in this chapter unchanged, i.e. the outperformance of macro-yields model does not carry over into the extended period.

Table 3.8: Out-of-sample forecast for excess bond returns in real time

	Vintage Data		Real-Time Data	
Maturity	1970:M1–2008:M12	1970:M1–2016:M12	1970:M1–2008:M12	1970:M1–2016:M12
2y	0.77	1.37	0.80	1.30
3y	0.75	1.30	0.78	1.26
4y	0.75	1.27	0.77	1.23
5y	0.76	1.28	0.79	1.24

Note. The table reports the rMSFE of the macro-yields model (MY) with respect to the expectations hypothesis for forecasts of 1-year-holding-period excess returns on 2- to 5-year bonds (in rows) over the extended timespan 1970:M1–2016:M12. The table shows the results using vintage (left panel) and real-time data (right panel), respectively.

3.4.3 Causes of Performance Deterioration

As discussed previously, the deterioration in model performance can be attributed to the excessive influence exerted by macroeconomic fluctuations through the vector autoregression system. To illustrate this, Figure 3.6 decomposes the 12-month-ahead forecast errors of the 1-year bond yield into those of the latent yield factors, and examines the generalised impulse response function (IRF) (Pesaran and Shin, 1998) of the latent five-variable VAR to examine how shocks to the macroeconomic panel reverberate through the system to influence the yield factors.

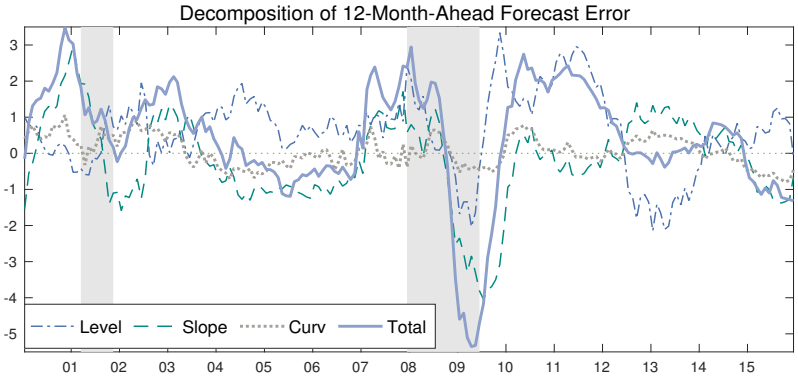

Figure 3.6: Decomposition of the 12-month-ahead forecast errors. The figure presents the decomposition of the 12-month-ahead forecast errors of 1-year bond yields into forecast errors of the underlying level, slope, and curvature factors, where the error is measured by $F_{\text{Model-implied}} - F_{\text{Realised}}$ with F={Yields, Level, Slope, Curvature}, over the timespan 2000:M1–2015:M12.

To enable exploration of the reasons behind the model performance, the 12-month-ahead forecast error of the 1-year bond yield is decomposed as shown in Figure 3.6. During the recession, the predictability decline is directly caused by the

under-prediction of bond yields, which can be further attributed to the simultaneous under-prediction of all three yield factors. All yield factor forecasts experience sharp declines during the recession, falling to unprecedented low levels. This disturbance, far exceeding the normal range of yield fluctuations, reflects the responses of the yield factors to the large exogenous macroeconomic shocks through the VAR system, as suggested by Figure 3.7. Of the three factors, the severe deviations of the yield slope forecasts from the actual evolvement (dashed line in Figure 3.7) account for the majority of the total yield forecast errors (solid line).

The tendency for the macroeconomic panel to revert to its normal level drives up the yield factors, most pronounced in the level factor (dash-dotted line), resulting in the considerable over-prediction of yields during the 2010–12 period. Consequently, the unspanned macro-yields model delivers inferior forecasts both during the financial meltdown and through the post-recession 2010–12 period, in line with the observed facts in Figure 3.4 and Figure 3.5.

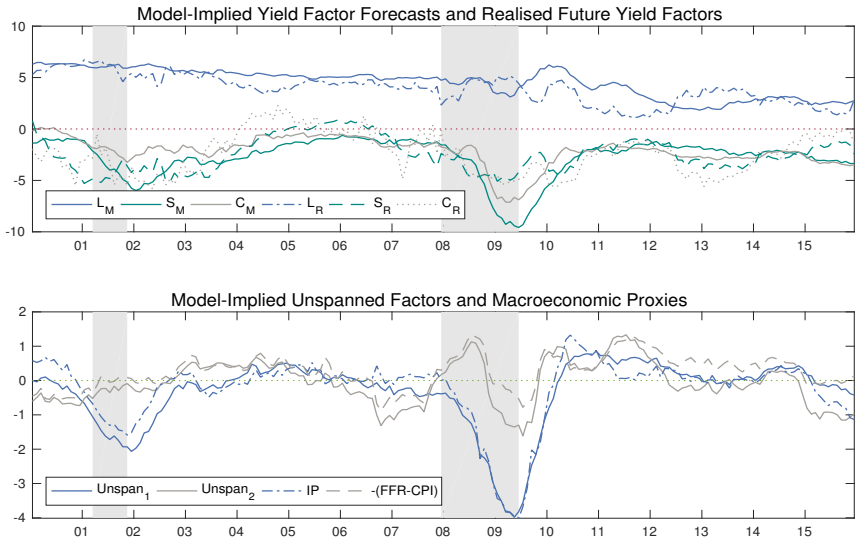

Figure 3.7: Model-implied and realised factors. The upper panel plots the model-implied yield factor forecasts (L_M, S_M, C_M) and realised future yield factors (L_R, S_R, C_R) over 2007:M1–2015:M12, with L/S/C referring to level/slope/curvature. The lower panel plots the model-implied unspanned factors ($Unspan_1$, $Unspan_2$) and their macroeconomic proxies, the industrial production (IP) growth rate, and the difference between the federal reserve rate (FFR) and consumer price index (CPI) growth rate.

Figure 3.8 reports the 12-month-ahead generalised impulse response functions of the yield factors in response to the innovations in the unspanned macro factors, where the impulse response functions are calculated using the parameters estimated recursively as per the out-of-sample forecasting analysis. The impulse response function tracks the future changes in the yield factors due to the one-unit

impulse in the current macro factors, and thus can be employed to explicitly measure how the model-implied forecasts respond to macroeconomic shocks.

As shown in Figure 3.8, the first (second) macro factors exhibit significant positive (negative) influences on the level and slope factors, which constitute the majority of the variations in yield movements. During the recession, the most notable negative shift occurs in the slope factor forecast. This deviation primarily stems from the impact exerted by the downward impulses in the first unspanned macro factor, one unit of which can lead to a 0.2 unit change in the slope factor one year into the future, according to the generalised impulse response function. Through this, the dramatic shocks in the macroeconomic panel exert their influence via the VAR equation on the yields.

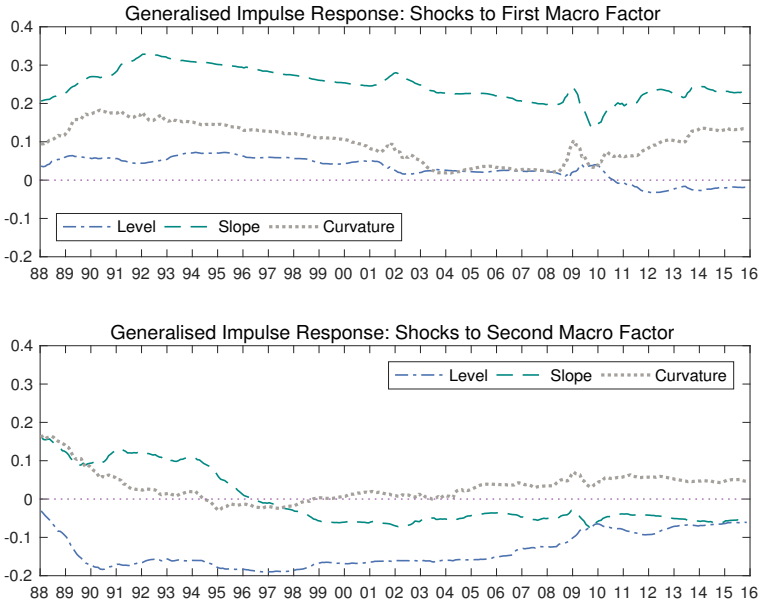

Figure 3.8: 12-month-ahead generalised impulse response: shocks to macro factors. The figure plots the generalised impulse responses of the three yield factors to the macro factor shocks at 12 months ahead. Based on the VAR parameters estimated recursively over extending windows, the impulse responses at 12 months ahead are calculated and reported in the figure, with the x-axis representing each forecast origin.

3.5 Stability during and after the Crisis

It has been argued that the decline in predictability in the aftermath of the crisis is primarily attributable to the short-term interest rates' oscillating in the constrained trajectory, which results in instabilities concentrated over the front-end of the curve, while the medium- and long-end remain relatively stable. To investigate this, the current section revisits the predictability after controlling for the

constrained movements of the front-end. This is achieved either by applying the conditional forecasting approach to construct counterfactual paths post-2008:M10, or by excluding the front-end completely from the yield construction to examine the differential effects. This section finds no evidence of substantially improved forecast accuracy, suggesting the possibility of instabilities in overall yield dynamics and macro-finance interactions in the context of macro-finance DFM modelling.

3.5.1 Counterfactual Assessment

Following the approach in Giannone, Lenza, and Reichlin (2019), this section explores counterfactual scenarios by conducting conditional forecast evaluation experiments, as outlined in Section 3.3.3. This method provides a useful avenue for examining the stability of the underlying predictive mechanisms, facilitating further insights into the yield dynamics and macro-financial linkages. Discrepancy between the counterfactual evolvement, which is compatible with the conditioning information, and the actual development would signal either instability of the dynamic interdependence or occurrence of unusual shocks. In this subsection, Figure 3.9 constructs in-sample conditional forecasts over the 1970:M1–2016:M12 period, where certain maturity spectra of the curve are regarded as unobservable after 2008:M10, as expressed in Equation (3.30), while Table 3.9 evaluates the out-of-sample forecast accuracy in this context.

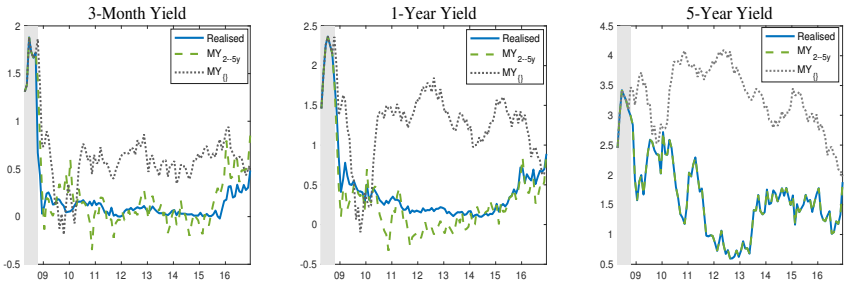

Figure 3.9: In-sample conditional forecasts of yields. This figure compares the realised paths (solid line) and the conditional forecasts of the 3-month (left), 1-year (middle), and 5-year (right) yields when certain maturity spectra of the curve are treated as unobservable after 2008:M10 and are projected by the observable variables (denoted by subscripts). In the "$\text{MY}_{\{2\text{–}5\text{y}\}}$" case (dashed line), the 3-month and 1-year yields are generated using the conditional forecasting approach while in an extreme scenario "$\text{MY}_{\{\}}$" (dotted line), the entire curve after 2008:M10 is re-constructed based on the pre-crisis yield dynamics and macro-yields relationships.

Excluding both the 3-month and 1-year yields from the conditioning set ($\text{MY}_{\{2\text{–}5\text{y}\}}$) generates projections that enter negative territory for both 3-month (left panel) and 1-year yields (middle panel). In an extreme case that recovers the entire yield curve based on the macroeconomic information, the macro-yields model

($MY_{\{\}}$) produces higher outcomes that depart substantially from the observed developments across all maturities, particularly from 2010 onwards, which is in line with the previous argument in Section 3.4 that the macro-yields model tends to over-predict in the recovery period. This discrepancy indicates that the post-crisis evolution of macroeconomic variables and the pre-crisis patterns together do not afford sufficient information to make accurate inferences across the full spectrum of rates, indicating changes in features of the yield distributions and/or in the dynamic relationship between macroeconomic and yield variables. To further gauge the impact of removing exogenous constraints on the short-term yields, and to assess the potential of the conditional forecasting approach as a possible solution for the post-crisis macro-finance modelling, Table 3.9 examines the out-of-sample forecast accuracy when certain maturity spectra of the curve are reconstructed post-2008:M10 conditional on the information set at each forecast origin.

Table 3.9: Counterfactual assessment in terms of yield curve forecasts

	3m	1y	2y	3y	4y	5y	3m	1y	2y	3y	4y	5y
			$MY_{\{1–5y\}}$						$MY_{\{2–5y\}}$			
h=1	1.22	1.12	1.12	1.12	1.09	1.07	1.68	1.63	1.14	1.12	1.08	1.07
h=3	1.19	1.22	1.23	1.21	1.20	1.18	1.36	1.32	1.24	1.21	1.19	1.18
h=6	1.26	1.29	1.34	1.35	1.33	1.31	1.33	1.32	1.35	1.35	1.33	1.31
h=12	1.21	1.23	1.34	1.43	1.46	1.44	1.23	1.24	1.35	1.43	1.46	1.43
h=24	1.09	1.14	1.36	1.56	1.68	1.81	1.10	1.14	1.36	1.56	1.68	1.81
			$MY_{\{3–5y\}}$						$MY_{\{4–5y\}}$			
h=1	2.68	2.19	1.29	1.12	1.09	1.08	3.83	3.64	1.92	1.25	1.12	1.12
h=3	1.56	1.41	1.26	1.22	1.22	1.20	1.81	1.62	1.35	1.26	1.28	1.28
h=6	1.38	1.35	1.37	1.38	1.37	1.35	1.43	1.39	1.43	1.47	1.49	1.50
h=12	1.25	1.27	1.39	1.49	1.53	1.51	1.30	1.35	1.54	1.69	1.76	1.76
h=24	1.17	1.22	1.46	1.68	1.81	1.94	1.40	1.48	1.79	2.06	2.22	2.39
			$MY_{\{5y\}}$						$MY_{\{\}}$			
h=1	3.94	3.87	2.08	1.31	1.14	1.12	3.13	7.47	12.92	16.23	17.50	18.82
h=3	1.83	1.64	1.36	1.26	1.27	1.28	1.88	2.77	4.15	5.21	5.84	6.31
h=6	1.44	1.39	1.42	1.46	1.49	1.50	1.72	2.13	2.86	3.45	3.75	3.94
h=12	1.30	1.35	1.54	1.70	1.77	1.77	1.61	1.82	2.29	2.70	2.94	3.05
h=24	1.42	1.50	1.82	2.09	2.25	2.42	1.61	1.74	2.14	2.50	2.73	2.98

Note. The table reports the rMSFE of the macro-yields model (MY) with respect to the random walk for forecasts of 3m-, 1y-, 2y-, 3y-, 4y-, and 5y-yields (in columns) made h=3, 6, 12, 24 months ahead (in rows) over the extended timespan 1970:M1–2016:M12. Each panel corresponds to one counterfactual scenario where certain maturity spectra of the curve are regarded as unobservable after 2008:M10 and re-constructed using the observable variables (denoted by subscripts).

Not surprisingly, the forecasts deteriorate as more yields are dropped from the set of conditioning variables, with the "$MY_{\{1–5y\}}$" case delivering the most accurate forecasts. Excluding the 3-month yield after 2008:M10 from the observation set appears to enhance the predictions of the short-end over medium- to long-term

horizons. This conditional forecasting attempt, however, does not make a noticeable difference for longer maturity bonds, leaving the average performance across the curve similar to that obtained in the original case $\text{MY}_{\{3\text{m}–5\text{y}\}}$. The no-predictability approach remains a challenging benchmark in this extended sample.

These findings also hold true for Table 3.10, which investigates the predictability of excess bond returns. The magnitude of differences across the OY scenarios, however, appears much smaller compared to that across the MY scenarios. In the original sample of 1970:M1–2008:M12, the MY model demonstrates substantial performance improvements compared with the OY model, with macroeconomic variables providing additional insights into the driving forces of yield factors. However, this predictive mechanism fails to operate effectively in terms of the post-crisis modelling, even after addressing the constrained movements of the front-end, suggesting the occurrence of instabilities during and after the crisis.

Table 3.10: Counterfactual assessment in terms of excess return forecasts

	$\text{MY}_{1970:2008}$	$\text{MY}_{\{3\text{m}–5\text{y}\}}$	$\text{MY}_{\{1–5\text{y}\}}$	$\text{MY}_{\{2–5\text{y}\}}$	$\text{MY}_{\{3–5\text{y}\}}$	$\text{MY}_{\{4–5\text{y}\}}$	$\text{MY}_{\{5\text{y}\}}$	$\text{MY}_{\{\}}$
2y	0.77	1.37	1.27	1.29	1.33	1.47	1.51	1.43
3y	0.75	1.30	1.26	1.25	1.27	1.41	1.45	1.30
4y	0.75	1.27	1.27	1.25	1.29	1.41	1.45	1.27
5y	0.76	1.28	1.29	1.27	1.32	1.48	1.48	1.25
	$\text{OY}_{1970:2008}$	$\text{OY}_{\{3\text{m}–5\text{y}\}}$	$\text{OY}_{\{1–5\text{y}\}}$	$\text{OY}_{\{2–5\text{y}\}}$	$\text{OY}_{\{3–5\text{y}\}}$	$\text{OY}_{\{4–5\text{y}\}}$	$\text{OY}_{\{5\text{y}\}}$	$\text{OY}_{\{\}}$
2y	1.20	1.23	1.29	1.28	1.25	1.25	1.31	1.25
3y	1.19	1.23	1.30	1.29	1.25	1.24	1.33	1.27
4y	1.18	1.18	1.26	1.25	1.20	1.21	1.33	1.27
5y	1.17	1.13	1.20	1.19	1.14	1.15	1.33	1.27

Note. The table reports the rMSFE of the macro-yields model (MY) with respect to the expectations hypothesis for forecasts of 1-year-holding-period excess returns on 2- to 5-year bonds (in rows) over the extended timespan 1970:M1–2016:M12. Each column corresponds to one counterfactual scenario where certain maturity spectra of the curve are regarded as unobservable after 2008:M10 and re-constructed using the observable variables (denoted by subscripts). For example, at each forecast origin t after 2008:M10, the conditioning information set of $\text{MY}_{\{2–5\text{y}\}}$ consists of 3m- and 1y-yields observations up to 2008:M10, and 2–5y yields and macroeconomic observations up to t. $\text{MY}_{1970:2008}$ and $\text{OY}_{1970:2008}$ refer to results of the original sample.

Collectively, both the in-sample and out-of-sample conditional forecast evaluations indicate considerable discrepancies between the actual evolvement of yields and those compatible with the post-crisis macroeconomic development. These results challenge the implicit assumptions of the conditional forecasting approach that the dependency structure and distribution of disturbances are consistent over time, implying instability in terms of the macro-financial linkages and yield dynamics.

3.5.2 "Ex-Front" vs "Full-Curve"

This subsection adopts a different perspective by comparing the "full-curve" scenario that includes all available yields with the "ex-front" scenario that excludes the short-end from the analysis, as displayed in Table 3.11 and Table 3.12. In the latter case, one might expect to observe a general decline in fit and forecasting performance owing to the omission of short-term rates, which correspond most closely with the monetary practice. Rather than focusing on absolute performance, this section develops conclusions based upon the differential effects before and after the crisis, examining the magnitude of changes with the inclusion of the post-global financial crisis sample. If the short-end behaviour is held responsible for the severe performance deterioration, then the post-crisis performance of the ex-front scenario should at least deteriorate to a much lesser extent relative to the full-curve scenario.

Table 3.11: Ex-front yield curve forecasts using Fama–Bliss dataset

	1970:M1–2008:M12					1970:M1–2016:M12				
$MY_{\{Full\text{-}Curve\}}$	1y	2y	3y	4y	5y	1y	2y	3y	4y	5y
h=1	0.92	1.00	1.00	0.99	0.99	1.11	1.12	1.10	1.07	1.06
h=3	0.92	0.98	0.96	0.98	0.98	1.24	1.23	1.19	1.17	1.16
h=6	0.89	0.94	0.93	0.93	0.94	1.35	1.37	1.34	1.31	1.29
h=12	0.75	0.79	0.80	0.81	0.81	1.33	1.39	1.43	1.44	1.41
h=24	0.65	0.72	0.78	0.83	0.90	1.19	1.35	1.51	1.62	1.74
$MY_{\{Ex\text{-}Front\}}$	1y	2y	3y	4y	5y	1y	2y	3y	4y	5y
h=1	0.91	0.98	0.99	0.97	0.96	1.07	1.08	1.07	1.02	1.02
h=3	0.96	1.02	0.99	0.99	1.00	1.25	1.22	1.16	1.13	1.13
h=6	0.96	1.01	0.99	0.98	1.00	1.37	1.36	1.32	1.29	1.30
h=12	0.83	0.90	0.91	0.92	0.94	1.35	1.39	1.42	1.44	1.45
h=24	0.77	0.82	0.88	0.93	1.01	1.35	1.49	1.64	1.76	1.92

Note. The table reports the rMSFE of the macro-yields model (MY) with respect to the random walk for the forecasts of 1y- to 5y-yields (in columns) made h=1, 3, 6, 12, 24 months ahead (in rows) using the Fama-Bliss dataset. The table shows the results of "Full-Curve" analysis (upper panel), which includes all available yields, and "Ex-Front" analysis (lower panel), which excludes the 3-month yield from the curve construction, over the original timespan 1970:M1–2008:M12 (left panel) and the extended timespan 1970:M1–2016:M12 (right panel).

As shown in Table 3.11, excluding the short-end from the analysis inevitably hampers predictability across the entire curve on both short and long time-scales. This aligns with the common notion that the monetary shock is most prominent on the short-term rates, the exclusion of which results in a material loss of monetary information. Despite the general decline in predictability, if the macro-yields model's predictive mechanisms remain valid over the longer end of the curve, the post-crisis performance should be comparable to that in the pre-crisis period. To say the least, the performance degradation over time would be noticeably lessened.

Controlling the possible exogenous constraints on the front-end, however, does not significantly alleviate the deterioration with the extended sample, as suggested by an average 0.38 worsening in the ex-front scenario, compared to 0.41 in the full-curve scenario. Similar patterns emerge for the excess bond return predictability presented in Table 3.12. Excluding the short-term rate from the analysis does not result in any distinguishable difference between the two scenarios, which also holds true for Le–Singleton yields that extend the spectrum up to 10 years (see Appendix 3.D).

Table 3.12: Ex-front excess return forecasts using Fama–Bliss dataset

	$MY_{\{Full\text{-}Curve\}}$				$MY_{\{Ex\text{-}Front\}}$			
	1970:M1–2008:M12		1970:M1–2016:M12		1970:M1–2008:M12		1970:M1–2016:M12	
Maturity	MY	OY	MY	OY	MY	OY	MY	OY
2y	0.77	1.20	1.37	1.23	0.86	1.18	1.40	1.21
3y	0.75	1.19	1.30	1.23	0.85	1.12	1.30	1.15
4y	0.75	1.18	1.27	1.18	0.85	1.09	1.26	1.09
5y	0.76	1.17	1.28	1.13	0.86	1.09	1.27	1.06

Note. The table reports the rMSFE of the macro-yields model (MY) and only-yields model (OY) with respect to the expectations hypothesis for forecasts of 1-year-holding-period excess returns on 2- to 5-year bonds (in rows) using the Fama-Bliss dataset. The table shows the results of "Full-Curve" analysis (left panel), which includes all available yields, and "Ex-Front" analysis (right panel), which excludes the 3m-yield from the curve construction, over the original timespan 1970:M1–2008:M12 and the extended timespan 1970:M1–2016:M12.

Table 3.11 and Table 3.12 imply that the deterioration in performance during and after the crisis cannot be ascribed solely to the short-term rate's continued low level. The dynamics and interdependence characterised by the macro-finance DFM do not appear to serve as a valid predictive mechanism in the extended sample, even after the removal of the front-end, suggesting the possibility that they might have altered in a manner that cannot be adequately accommodated by the DFM.

3.6 Conclusion

While the in-sample results exhibit moderate changes, the out-of-sample forecasting performance deteriorates substantially, leading to different inferences regarding the macro-yields model. During the recession and recovery periods, evidence in favour of the five-factor model weakens due to the shift in the dependence structure assumed by the linear-Gaussian DFM. The significant decline in the predictive power can be partly attributed to the inferior performance arising from the Gaussian specification during and after the crisis. Directions for future research include time-varying parameter models that allow for changes in the dependence structure,

and non-Gaussian state-space models that allow for skewed and heavy-tailed innovations in the measurement and transition equations. To address the latter, Chapter 4 develops an effective modelling strategy that potentially deals better than the linear-Gaussian DFM with the post-2008 sample.

3.7 Appendices

3.A Data Source

Macroeconomic variables include average hourly earnings (AHE), consumer price index (CPI), real disposable personal income (INC), effective federal funds rate (FFR), new one family houses sold (HSal), industrial production index (IP), M1 money stock (M1), ISM manufacturing PMI composite index (Manf), total nonfarm payrolls (Paym), personal consumption expenditures (PCE), producer price index: crude materials for further processing (PPIc), producer price index: finished goods (PPIf), capacity utilisation (CU), civilian unemployment (Unem). Except FFR, Manf, CU, and Unem, annual growth rates of the other variables are obtained to ensure stationarity, according to Coroneo, Giannone, and Modugno (2016). This chapter collects data from the Federal Reserve Economic Data (FRED) and Bloomberg.

3.B Diebold–Mariano Test

This section reports the test proposed by Diebold and Mariano (2002) of the forecast differences examined in Table 3.5 and Table 3.6 on the basis of the squared-error loss differential $d = e^2_{\text{Model}} - e^2_{\text{RW}}$. To test the null hypothesis of equal predictive accuracy at h steps ahead, i.e. $H_0 : E[d_t] = 0$, the large-sample $N(0,1)$ Diebold–Mariano statistic is computed as $DM = \frac{\bar{d}}{\sqrt{2\pi \hat{f}_d(0)/T}}$, where $\bar{d}$ denotes the sample average. To take account of the autocorrelation in h-step-ahead forecast errors, the variance of $\bar{d}$ is estimated via

$$2\pi \hat{f}_d(0) = \sum_{\tau=-(T-1)}^{(T-1)} \omega_\tau \gamma_d(\tau) \tag{3.31}$$

while the sample autocovariances $\gamma_d(\tau)$ for time displacement τ are obtained by $\frac{1}{T}\sum_{t=|\tau|+1}^{T}(d_t - \bar{d})(d_{t-|\tau|} - \bar{d})$. When correcting for serial correlation, the Newey-West heteroscedasticity and autocorrelation consistent estimator utilises a triangular window that assigns weights $\omega_\tau = \frac{h-\tau}{h}$ to autocovariances $\gamma_d(\tau)$ for $\tau < h$.

Table 3.13 and Table 3.14 report the above DM statistics for equal predictive ability of the model-based forecasts and the no-predictability benchmarks, where negative represents reduction in relative forecast error with respect to the random walk (Table 3.13) and the expectations hypothesis (Table 3.14). Against the two-sided alternative at significance level $\alpha = 10\%$, the null hypothesis of equal accuracy is rejected when DM falls outside [-1.645, 1.645]. In pre-2009 sample, the macro-yields model demonstrates accuracy improvement for shorter maturities at longer horizons as well as comparable performance in other cases, according to the sign and magnitude of DM statistics. Over the extended period, the macro-yields model substantially underperforms the random walk at shorter horizons while the difference is substantially discounted at the longest horizon (24-step-ahead) due to cumulation of the autocovariance terms. Examining the predictability of excess bond returns, Table 3.14 reveals a similar pattern that while the macro-yields approach delivers significantly enhanced forecasting before the recession, it results in declined but indistinguishable forecast accuracy relative to the expectations hypothesis in the full sample that ends in 2016.

Table 3.13: Diebold–Mariano statistics for out-of-sample yield curve forecasts

	1970:M1 – 2008:M12						1970:M1 – 2016:M12					
	Macro-Yields											
Maturity	3m	1y	2y	3y	4y	5y	3m	1y	2y	3y	4y	5y
h=1	**-1.77**	-0.29	0.55	0.43	0.37	0.16	**2.45**	**3.64**	**3.46**	**2.27**	**2.66**	**3.38**
h=3	-0.76	0.46	0.64	0.38	0.62	0.69	1.49	**2.07**	**2.35**	**2.17**	**2.20**	**2.33**
h=6	-0.34	0.37	0.49	0.36	0.38	0.41	1.44	**1.76**	**2.06**	**2.10**	**2.10**	**2.09**
h=12	-1.11	-0.70	-0.60	-0.57	-0.51	-0.54	0.97	1.14	1.44	1.60	**1.68**	**1.65**
h=24	**-2.25**	**-2.19**	**-1.96**	**-1.66**	-1.15	-0.43	0.31	0.37	0.72	1.00	1.16	1.33
	Only-Yields											
Maturity	3m	1y	2y	3y	4y	5y	3m	1y	2y	3y	4y	5y
h=1	**-2.94**	1.37	**3.62**	**2.45**	**2.26**	**2.11**	**-2.56**	**1.90**	**3.70**	**2.89**	**2.07**	**2.69**
h=3	-1.29	**1.97**	**2.83**	**2.08**	**2.02**	**2.01**	-0.81	**1.69**	**3.03**	**2.49**	**2.10**	**2.19**
h=6	-0.59	**1.77**	**2.43**	**2.19**	**1.91**	**1.90**	-0.02	**1.87**	**2.76**	**2.61**	**2.08**	**1.98**
h=12	0.05	1.03	1.55	1.63	1.60	1.63	0.47	1.31	**1.93**	**2.06**	**1.89**	**1.75**
h=24	0.31	0.66	1.30	**1.75**	**2.06**	**2.42**	0.90	1.22	**1.92**	**2.30**	**2.40**	**2.51**

Note. The table reports the Diebold–Mariano test statistics to test the null hypothesis of no difference between the random walk benchmark and the macro-yields(upper)/only-yields(lower) model in terms of forecasts of 3m-, 1y-, 2y-, 3y-, 4y-, and 5y-yields (in columns) made h=1, 3, 6, 12, 24 months ahead (in rows). The left panel shows the pre-2009 results, and the right panel displays the results of the extended sample. Against the two-sided alternative at significance level $\alpha = 10\%$, the null hypothesis of equal accuracy is rejected when DM falls outside [-1.645, 1.645], as marked in bold.

3.C Performance over Original Sample

In this section, the in-sample and out-of-sample analyses are carried out over the original timespan to check the validity of the estimation and the computation. The

Table 3.14: Diebold–Mariano statistics for out-of-sample bond return forecasts

	1970:M1 – 2008:M12					1970:M1 – 2016:M12				
Maturity	MY	OY	CP	LN	LN+CP	MY	OY	CP	LN	LN+CP
2y	-1.59	0.97	0.72	-0.94	-1.07	1.01	1.15	1.31	0.75	-0.22
3y	**-1.95**	1.04	0.88	-1.05	-0.89	0.98	1.25	1.60	0.31	-0.50
4y	**-2.05**	0.94	0.87	-1.11	-0.85	0.97	1.03	**1.69**	-0.16	-0.84
5y	**-2.05**	0.92	0.83	-1.07	-0.89	1.05	0.77	**1.74**	-0.53	-1.06

Note. The table reports the Diebold–Mariano test statistics to test the null hypothesis of no difference between the expectations hypothesis benchmark and different approaches in terms of forecasts of 1-year-holding-period excess returns on 2- to 5-year bonds (in rows). The columns correspond to the macro-yields model (MY), the only-yields model (OY), the CP factor, the LN factor, and the joint LN+CP factor, respectively. The left panel shows the pre-2009 results, and the right panel displays the results of the extended sample. Against the two-sided alternative at significance level $\alpha = 10\%$, the null hypothesis of equal accuracy is rejected when DM falls outside [-1.645, 1.645], as marked in bold.

narrow replication almost completely reproduces the results of Coroneo, Giannone, and Modugno (2016) over the 1970:M1–2008:M12 period. In-sample results, including the information criteria for model selection displayed in Table 3.15, the cumulative variance explained by the macro-yields factors in Table 3.16, the goodness-of-fit for the excess bond returns in Table 3.17, and the likelihood ratio test statistic for the unspanning restrictions in Table 3.18, are virtually identical to the original results (± 0.01) over the original sample.

Table 3.15: Selecting the number of macro factors

	Original		Replication	
Number of factors	IC	V	IC	V
3	0.02	0.44	0.02	0.44
4	-0.03	0.31	-0.03	0.31
5	-0.11	0.22	-0.11	0.22
6	0.01	0.18	0.01	0.18
7	0.23	0.17	0.23	0.17
8	0.43	0.16	0.42	0.16

Note. The table reports the modified information criterion (IC) over the original timespan 1970:M1–2008:M12. Results are reported for k-factor (k=3,4,5,6,7,8) models in rows, with the left panel corresponding to the original results reported by Coroneo, Giannone, and Modugno (2016) and the right panel to the replication results.

For the out-of-sample results, there are only tiny differences in the only-yields model between the replication results and the original results. The replication differences for the macro-yields model are slightly larger, which might be attributed to the increased computational complexity introduced by the additional macro parameters. With the slight variations reflecting tiny differences in the choice of the evaluation period, the replication results on the relative MSFE for both models in Table 3.19 are close to the original results, except for the 1-month-ahead forecasts. In Table 3.20, the replication concerning the 12-month-ahead predictive performance for the 1-year-holding-period excess returns provides similar results (± 0.01). Ac-

Table 3.16: Cumulative variance explained by the macro-yields factors

	Original					Replication				
	Level	Slope	Curv	UM1	UM2	Level	Slope	Curv	UM1	UM2
3m	0.59	0.94	1.00	1.00	1.00	0.59	0.94	1.00	1.00	1.00
1y	0.61	0.83	0.99	0.99	0.99	0.61	0.83	0.99	0.99	0.99
2y	0.65	0.78	0.99	0.99	0.99	0.65	0.78	0.99	0.99	0.99
3y	0.70	0.79	1.00	1.00	1.00	0.70	0.79	1.00	1.00	1.00
4y	0.74	0.80	0.99	0.99	0.99	0.74	0.80	0.99	0.99	0.99
5y	0.78	0.82	0.99	0.99	0.99	0.78	0.82	0.99	0.99	0.99
AHE	0.07	0.29	0.33	0.33	0.67	0.07	0.29	0.33	0.33	0.67
CPI	0.19	0.48	0.48	0.50	0.85	0.19	0.49	0.48	0.50	0.85
INC	0.00	0.02	0.03	0.34	0.36	0.00	0.02	0.03	0.34	0.36
FFR	0.53	0.93	0.96	0.96	0.97	0.53	0.93	0.96	0.96	0.97
Hsal	0.00	0.19	0.19	0.23	0.23	0.00	0.20	0.19	0.23	0.23
IP	0.02	0.02	0.03	0.69	0.69	0.02	0.02	0.03	0.70	0.70
M1	0.17	0.25	0.25	0.25	0.31	0.17	0.25	0.26	0.25	0.31
Manf	0.03	0.05	0.05	0.61	0.65	0.03	0.06	0.05	0.61	0.65
Paym	0.00	0.02	0.10	0.70	0.70	0.00	0.02	0.10	0.71	0.70
PCE	0.16	0.23	0.33	0.46	0.78	0.16	0.23	0.33	0.46	0.78
PPIc	0.03	0.14	0.14	0.20	0.43	0.03	0.14	0.14	0.20	0.43
PPIf	0.03	0.32	0.32	0.33	0.80	0.03	0.34	0.32	0.33	0.80
CU	0.02	0.16	0.21	0.63	0.64	0.02	0.16	0.21	0.63	0.64
Unem	0.44	0.54	0.55	0.65	0.68	0.44	0.54	0.55	0.65	0.68

Note. In each row, the table reports the cumulative variance of each yield or macro variable explained by the yield curve factors (Level, Slope, and Curv) and the unspanned macro factors (UM1 and UM2) over the original timespan 1970:M1–2008:M12, with the left panel corresponding to the original results and the right panel to the replication results.

Table 3.17: In-sample predictability of excess bond returns

	Original					Replication				
Maturity	MY	OY	CP	LN	LN+CP	MY	OY	CP	LN	LN+CP
2y	0.55	0.12	0.22	0.33	0.41	0.55	0.12	0.22	0.33	0.41
3y	0.53	0.12	0.24	0.33	0.43	0.53	0.12	0.24	0.33	0.42
4y	0.50	0.14	0.27	0.32	0.43	0.49	0.14	0.27	0.32	0.43
5y	0.46	0.15	0.24	0.30	0.40	0.45	0.15	0.24	0.30	0.40

Note. This table reports the proportion of the variation in 1-year-ahead 1-year-holding-period excess returns accounted for by the macro-yields model (MY), the only-yields model (OY), the CP factor (variables constructed using 12-, 24-, 36-, 48-, 60-month forward rates), the LN factor (combinations of principal components extracted from 131 macroeconomic data series), and the joint LN+CP factor. The rows correspond to excess returns on 2- to 5-year bonds. The left panel shows the original results, and the right panel displays the replication results.

Table 3.18: Likelihood ratio test for the unspanning restrictions

	Original		Replication	
H_0	Test statistic	p-value	Test statistic	p-value
$\Gamma_{yx} = 0$	12.85	0.38	12.85	0.38
$A_{yx} = 0$	79.03	0.00	79.03	0.00

Note. The table reports the likelihood ratio test statistics and p-values for the null hypotheses $\Gamma_{yx} = 0$ in Equation (3.9) and $A_{yx} = 0$ in Equation (3.10) over the original timespan 1970:M1–2008:M12, with the left panel corresponding to the original results and the right panel to the replication results.

cordingly, the main findings and conclusions that suggest the outperformance of the macro-yields model still hold.

Table 3.19: Out-of-sample predictability of the yield curve

	Original						Replication					
	Macro-yields											
Maturity	3m	1y	2y	3y	4y	5y	3m	1y	2y	3y	4y	5y
h=3	0.79	0.93	0.99	0.96	0.99	1.02	0.79	0.92	0.98	0.96	0.98	0.98
h=6	0.78	0.89	0.94	0.93	0.93	0.94	0.79	0.89	0.94	0.93	0.93	0.94
h=12	0.69	0.74	0.79	0.80	0.80	0.80	0.69	0.75	0.79	0.80	0.81	0.81
h=24	0.62	0.66	0.74	0.82	0.88	0.97	0.61	0.65	0.72	0.78	0.83	0.90
	Only-yields											
Maturity	3m	1y	2y	3y	4y	5y	3m	1y	2y	3y	4y	5y
h=3	0.96	1.13	1.20	1.14	1.10	1.13	0.96	1.14	1.21	1.14	1.11	1.11
h=6	0.99	1.18	1.25	1.21	1.15	1.16	0.99	1.18	1.26	1.22	1.17	1.17
h=12	1.04	1.16	1.26	1.27	1.25	1.26	1.05	1.16	1.26	1.27	1.25	1.25
h=24	1.06	1.12	1.27	1.39	1.49	1.62	1.09	1.14	1.28	1.40	1.49	1.61

Note. The table reports the rMSFE of the macro-yields model (upper panel) and the only-yields model (lower panel) with respect to the random walk for forecasts of 3m-, 1y-, 2y-, 3y-, 4y-, and 5y-yields (in columns) made h=1, 3, 6, 12, 24 months ahead (in rows). The left panel shows the original results, and the right panel displays the replication results over the original timespan 1970:M1–2008:M12.

Table 3.20: Out-of-sample predictability of excess bond returns

	Original					Replication				
Maturity	MY	OY	CP	LN	LN+CP	MY	OY	CP	LN	LN+CP
2y	0.76	1.20	1.17	0.80	0.80	0.77	1.20	1.17	0.79	0.80
3y	0.75	1.20	1.21	0.79	0.83	0.75	1.19	1.20	0.78	0.83
4y	0.74	1.18	1.21	0.78	0.83	0.75	1.18	1.20	0.78	0.83
5y	0.75	1.18	1.18	0.81	0.83	0.76	1.17	1.18	0.80	0.83

Note. The table reports the rMSFE of different approaches with respect to the expectations hypothesis for forecasts of 1-year-holding-period excess returns on 2- to 5-year bonds (in rows). The columns correspond to the macro-yields model (MY), the only-yields model (OY), the CP factor, the LN factor, and the joint LN+CP factor, respectively. The left panel shows the original results, and the right panel displays the replication results over the original timespan 1970:M1–2008:M12.

3.D Analysis using Le–Singleton Yields

One might wonder whether the results in Section 3.5.2 may be partly due to the limited number of observations after dropping the short-term ones, over which the full potential of the Nelson–Siegel extrapolation cannot be unleashed. To examine this issue, Table 3.21 revisits the two scenarios using the yield data of Le and Singleton (2013) that extends the spectrum up to 10 years. Despite mixed results regarding the absolute performance, the ex-front scenario does not substantially mitigate the deterioration, as indicated by the 0.35 increase in rMSFE for the ex-front scenario and 0.39 for the full-curve scenario. The macro-yields model loses its appeal against the random walk, generating an average rMSFE of 1.30 in the

full sample compared to 0.95 in the pre-crisis sample. Predictability regarding the excess bond returns as reported in Table 3.22 collaborates these findings, as shown by the 0.41 increase in rMSFE for the ex-front scenario and 0.45 for the full-curve scenario.

Table 3.21: Ex-front yield curve forecasts using Le–Singleton yields

	$MY_{\{Full\text{-}Curve\}}$										$MY_{\{Ex\text{-}Front\}}$									
	1970:M1–2008:M12																			
	1y	2y	3y	4y	5y	6y	7y	8y	9y	10y	1y	2y	3y	4y	5y	6y	7y	8y	9y	10y
1m	1.03	1.10	1.09	1.06	1.04	1.04	1.00	1.01	1.01	0.98	0.95	1.04	1.04	1.02	1.01	1.01	0.98	1.00	1.00	0.98
3m	1.01	1.09	1.06	1.05	1.04	1.05	1.03	1.02	1.02	1.01	0.92	1.03	1.01	1.02	1.01	1.02	1.02	1.03	1.02	1.04
6m	0.89	1.00	1.00	1.01	1.01	1.02	1.01	1.01	1.01	1.01	0.82	0.94	0.96	0.98	0.98	1.00	1.01	1.02	1.02	1.04
12m	0.67	0.77	0.81	0.84	0.86	0.89	0.90	0.92	0.93	0.92	0.66	0.77	0.82	0.86	0.88	0.91	0.93	0.95	0.96	0.98
24m	0.59	0.62	0.66	0.69	0.74	0.78	0.81	0.83	0.84	0.85	0.70	0.73	0.78	0.82	0.88	0.92	0.95	0.98	0.99	0.99
	1970:M1–2016:M12																			
1m	1.16	1.18	1.16	1.10	1.09	1.08	1.04	1.06	1.05	1.03	1.12	1.14	1.12	1.07	1.06	1.06	1.03	1.05	1.05	1.04
3m	1.24	1.27	1.22	1.20	1.18	1.19	1.15	1.16	1.14	1.14	1.23	1.24	1.19	1.17	1.16	1.18	1.16	1.18	1.17	1.20
6m	1.27	1.33	1.32	1.33	1.32	1.34	1.31	1.31	1.28	1.29	1.28	1.31	1.29	1.30	1.30	1.34	1.32	1.34	1.33	1.37
12m	1.19	1.26	1.33	1.39	1.43	1.46	1.46	1.47	1.48	1.47	1.18	1.23	1.29	1.36	1.40	1.43	1.45	1.48	1.49	1.51
24m	1.14	1.23	1.36	1.47	1.62	1.73	1.80	1.90	1.93	1.99	1.05	1.14	1.28	1.39	1.54	1.64	1.71	1.79	1.82	1.86

Note. The table reports the rMSFE of the macro-yields model (MY) with respect to the random walk for the forecasts of 1y- to 10y-yields (in columns) made h=1, 3, 6, 12, 24 months ahead (in rows) using the Le–Singleton yields. The table shows the results of "Full-Curve" analysis (left panel), which includes all available yields, and "Ex-Front" analysis (right panel), which excludes the 3m-yield from the curve construction, over the original timespan 1970:M1–2008:M12 (upper) and the extended timespan 1970:M1–2016:M12 (lower).

Table 3.22: Ex-front excess return forecasts using Le–Singleton yields

	$MY_{\{Full\text{-}Curve\}}$				$MY_{\{Ex\text{-}Front\}}$			
	1970:M1–2008:M12		1970:M1–2016:M12		1970:M1–2008:M12		1970:M1–2016:M12	
	MY	OY	MY	OY	MY	OY	MY	OY
2y	0.68	1.19	1.22	1.27	0.67	1.16	1.21	1.20
3y	0.74	1.15	1.20	1.24	0.74	1.13	1.16	1.18
4y	0.75	1.10	1.18	1.18	0.76	1.09	1.15	1.12
5y	0.79	1.09	1.23	1.16	0.80	1.09	1.20	1.10
6y	0.81	1.11	1.27	1.16	0.83	1.12	1.24	1.10
7y	0.79	1.06	1.22	1.06	0.81	1.07	1.20	1.02
8y	0.82	1.06	1.26	1.05	0.84	1.08	1.25	1.01
9y	0.85	1.08	1.28	1.03	0.88	1.11	1.29	1.01
10y	0.81	1.04	1.22	0.97	0.84	1.07	1.22	0.95

Note. The table reports the rMSFE of the "Full-Curve" and "Ex-Front" scenarios with respect to the expectations hypothesis for forecasts of 1-year-holding-period excess returns on 2- to 10-year bonds (in rows) using the Le–Singleton yields, over the original timespan 1970:M1–2008:M12 and the extended timespan 1970:M1–2016:M12.

Chapter 4

Macro-Yields Modelling in the Presence of Asymmetrically Distributed Interest Rates

4.1 Introduction

Until recently, there has been limited interest in investigating and exploiting the non-Gaussianity in yield distributions when modelling and forecasting government bonds. The vast majority of dynamic term structure models, with or without a macroeconomic flavour, can be regarded as belonging to the class of linear-Gaussian models, which admit tractable pricing implications and efficient estimation procedures. These approaches typically assume a standard Gaussian VAR process featuring symmetrically-distributed innovations for the underlying states that drive the economy. Accordingly, it generates margins that follow Gaussian or a mixture of Gaussian distributions, regardless of whether the arbitrage-free affine framework or the Nelson–Siegel representation is adopted to characterise the term structure. The validity of this assumption has not received considerable critical attention until recently, while it is becoming increasingly difficult to ignore the existence of asymmetry in the distribution of bond yields, in particular for short maturities. Drawing upon the insights of Smith and Vahey (2016), which highlight the importance of taking into account asymmetry properties when forecasting macroeconomic variables, this research aims to advance the understanding of the non-Gaussian aspects of government bond yields, and to explore their implications in the context of macro-finance term structure modelling through implementation of copula methodologies.

Exploiting the state-space inversion copula, the non-Gaussian macro-yields model proposed here allows for asymmetric, multimodal, and heavy-tailed marginal distributions, while retaining the dynamic factor structure that explicitly addresses cross-sectional and serial dependence, in order to accommodate both predictive causality and contemporaneous correlations. The model aims to borrow the de-

pendence structure that underlies the macro-yields model of Coroneo, Giannone, and Modugno (2016), which effectively characterises the relationships among a relatively large collection of macro-yields variables in a parsimonious manner, while relaxing the commonly held assumption of Gaussianity via fitting asymmetrical non-parametric margins. This can be achieved by extracting the implicit copula that inherits only the dependencies implied by the macro-yields model, which has the potential to generate more accurate forecasts when combined with appropriate choices of marginal distributions. In doing so, the proposed model can be regarded as a generalisation of the original macro-yields model, encompassing the latter as a special case that deploys normal distributions for all margins, while promoting greater flexibility by plausibly accommodating the univariate features.

In terms of out-of-sample forecast accuracy over the entire 1970:M1–2016:M12 period, the non-Gaussian macro-yields model outperforms its Gaussian counterpart across the maturity spectrum and forecast horizons examined. Inducing non-Gaussianity via copulas leads to an average 0.13 reduction in relative mean squared forecast errors (rMSFE) with respect to the random walk for the only-yields term structure model and an even larger 0.39 reduction for the macro-yields model, with notable decreases in multi-step-ahead forecasts errors of short-term yields. In comparison with the no-predictability benchmark, the proposed method demonstrates enhanced forecasting capability along the entire yield curve, delivering improved medium- to long-term predictions and comparable short-term forecasts.

A further examination of the excess return predictability provides more convincing evidence that supports the adoption of the non-Gaussian setup employing copula-based techniques. Utilising non-Gaussian schemes significantly adds to the in-sample and out-of-sample predictability of excess bond returns, resulting in an average 0.15 reduction in rMSFE with respect to the expectations hypothesis for the only-yields model and a substantial 0.59 reduction for the macro-yields model. Furthermore, it outperforms the predictive regressions with the most promising predictors by 8.2% in terms of R^2 and by 0.16 in terms of rMSFE. Overall, the non-Gaussian macro-yields model affords more accurate forecasts relative to its Gaussian counterpart and other prevailing forecasting approaches. Closer inspection reveals that, contrary to expectations, the inclusion of macroeconomic variables into the term structure modelling impairs its ability to forecast within the conventional Gaussian framework. Only by means of copulas that accommodate non-Gaussian behaviour does the macroeconomic information contribute to forecast accuracy over this extended period of time.

The enhancement enabled by the proposed model can be largely attributed to its satisfactory performance in the post-global financial crisis period when the non-

Gaussian effects become more pronounced. The non-Gaussian macro-yields model adequately captures yields' inclination towards the positive end at or near zero, and plausibly adapts the marginal distributions to account for the emerging multimodality during the post-recession period, thus producing more reasonable forecasts. Its Gaussian rival, however, assigns higher probabilities to negative forecasts during the crisis and is constrained to be unimodal, resulting in severe overprediction after the recession. Examining the performance through time further assesses the stability and robustness of the copula-based approach. The non-Gaussian scheme enables enhanced predictions for the majority of the time with respect to its Gaussian counterpart and the no-predictability benchmark, neither specific to a particular time-period nor subject to a one-time reversal. The improvement is most prominent in periods following recessions, such as the first half of the 1990s and the economic expansion of 2002–2007, and also holds true in recent years.

This chapter has gone some way to furthering our understanding of the predictive mechanisms underlying the macro-finance modelling, and presenting possible solutions to the challenges posed by the recent recession. This study draws upon the strand of research that deals with serial as well as cross-sectional dependence in low-dimensional macroeconomic data using implicit copulas (Smith and Vahey, 2016; Smith and Maneesoonthorn, 2018; Loaiza-Maya and Smith, 2019). By additionally imposing a dynamic factor structure for dimensionality reduction, this study represents a productive attempt to explore the usefulness of the state-space inversion copula formulation in dealing with high-dimensional datasets in the macro-finance context. In doing so, this study provides an empirical avenue for exploring the potential of the copulas in the vast field of macro-finance research in a data-rich environment. The implicit copula approach explored here can conveniently retrieve inherent structures from any established multivariate representations without being limited to existing canonical distributions, contributing towards a more comprehensive characterisation of the interactions between the macroeconomy and asset markets.

The rest of the paper proceeds as follows. Section 4.2 provides a brief review of the relevant literature, while Section 4.3 describes the formulation and implementation of the non-Gaussian macro-yields model, which is constructed from the state-space inversion copula in the context of macro-finance term structure modelling. Section 4.4 examines the in-sample and out-of-sample predictability, before Section 4.5 discusses other implications while Section 4.6 concludes and speculates on future directions.

4.2 Related Literature

In the case of macro-finance modelling, empirical research has revealed causal linkages between macroeconomic fundamentals and the yield curve, indicating that the macro variables contribute insights into future yield movements. According to many in the field, macroeconomic variables enhance the yield forecasting by providing additional information regarding the real economic activity and inflation. Specific macroeconomic variables, such as the GDP growth and inflation (Wright, 2011) or the Chicago Fed National Activity Index and surveyed inflation forecasts (Joslin et al., 2014), are selected as representative indicators and included directly in the factor set. Holding the view that exact measures of aggregate activity and inflation are unobtainable, some studies perform dimension reduction procedures, such as PCA and factor analysis (Stock and Watson, 2002; Forni et al., 2005) to construct summary indicators from a group of variables (Ang and Piazzesi, 2003). Research attempts to investigate the "data-rich environment", referring to a broader panel consisting of tens to hundreds of macroeconomic measures, facilitated by advances in factor analysis (Bernanke and Boivin, 2003; Giannone et al., 2004; Bernanke et al., 2005; Favero et al., 2005; Mönch, 2008; Favero et al., 2012).

The term "unspanned" has been utilised to imply that the macroeconomic components contain predictive information beyond that afforded by the yield curve itself. In this vein, Coroneo, Giannone, and Modugno (2016) formulate the unspanned macro-yields model that extracts yield information from a relatively large panel of macroeconomic indicators. Chapter 3 in this thesis finds that, despite improving the accuracy of yield fitting and forecasting within the 1970:M1–2008:M12 period, this macro-yields model experiences severe performance deterioration in an extended sample that spans 1970:M1 to 2016:M12. Viewed as an application of the dynamic factor approach, the macro-yields model complies with underlying assumptions of Gaussian innovations and linear dependence, and thus is liable to suffer substantial performance degradation due to its limited ability to characterise skewed and heavy-tailed margins. This effect is especially pronounced after the onset of the 2007–09 turmoil, since when the yield curve exhibits more notable deviations from Gaussian behaviour. This chapter adds to this vein of research with the intention of extending this macro-yields model by relaxing the common but often unrealistic assumption of Gaussianity, in order to generate more plausible forecasts in the full 1970:M1–2016:M12 sample that covers the crisis and the post-crisis period.

Another strand of relevant literature focuses on a popular method for multivariate dependence characterisation, copula modelling, which is a widely-recognised approach to depicting various joint distributions that include but are not limited to

existing multivariate representations. Although extensive research has been carried out regarding its performance in economic forecasting (Teräsvirta, 2006; Patton, 2012, 2013; Smith and Vahey, 2016) and financial modelling related to credit risk and financial crisis (Li, 2000; Andersen and Sidenius, 2004; Laurent and Gregory, 2005; Hull and White, 2006; Christoffersen et al., 2012; Lucas et al., 2014; Creal and Tsay, 2015; Christoffersen et al., 2017; Oh and Patton, 2017, 2018), there is a paucity of studies that investigate its applications in the interdisciplinary macro-finance domain. Moreover, in terms of yield modelling, barely any attention has been paid to the non-linear term structure dependence to date, except the single work by Junker, Szimayer, and Wagner (2006) that breaks with tradition by questioning the linearity in the pure-yields models. Although the non-Gaussian macro-yields model proposed here employs an underlying linear dependence for the latent process, it does not imply a linear dependence structure for the observed variables, after taking into account the effects of non-Gaussian marginal distributions. Moreover, this copula-based approach is readily extensible to explicitly accommodate non-linearity by adopting non-linear state-space inversion copulas, as illustrated in Smith and Maneesoonthorn (2018) and Loaiza-Maya and Smith (2019).

The flexibility afforded by copula modelling stems from its broad range of choices in terms of marginal distributions and copulas. The Sklar's theorem lays the theoretical foundation for combining arbitrary marginal distributions through various dependence structures to construct valid joint distributions. On the question of macro-finance modelling, one major issue concerns the large number of macroeconomic and yield variables involved, from a condensed set of five factors to dozens or even hundreds. Recent advances in copula methods have facilitated investigation of high-dimensional dependence, making inference and estimation of the large collection feasible. These copula models address higher-dimensional dependence by enhancing the covariance matrix estimation (Engle and Kelly, 2012; Fan et al., 2012; Hautsch et al., 2012), employing vine-copulas (Aas et al., 2009; Min and Czado, 2010; Smith, 2015; Almeida et al., 2016), and adopting a factor modelling approach (Andersen and Sidenius, 2004; Laurent and Gregory, 2005; Krupskii and Joe, 2013; Creal and Tsay, 2015; Oh and Patton, 2017, 2018).

Besides the additional flexibility, the copula approach establishes a quantitative framework for examining various types of dependency such as tail dependence, heterogeneous dependence, and asymmetric dependence inherent in the variables, which might indicate directions for future research. The methodology applied in this research is related to that considered by Patton and colleagues (Patton, 2012; Oh and Patton, 2017, 2018), but a crucial difference lies in how serial dependences are taken account of. From a formulation perspective, models intended to accom-

modate causality also need to be capable of explicitly addressing the cross-serial dependency. The state-space inversion copula employed here enables a considerable amount of cross-serial as well as cross-sectional dependence by inheriting the dependence structures that underlie the state-space models, along with flexible choices of marginal distributions that reproduce univariate stylised facts.

Thanks to the copula-marginal separation and practical implementation of state-space models, the attractive features outlined above do not necessarily come at the cost of increased complexity in formulation and estimation processes. Through the inversion method, the dependence structure exploited is essentially the implicit copula drawn from a DFM (Coroneo et al., 2016; Stock and Watson, 2016) that summarises the co-movements of a large cross section of variables into a reduced number of autoregressive factors while allowing some degree of serial correlation in idiosyncratic disturbances. The canonical maximum likelihood estimation (Genest et al., 1995; Tsukahara, 2005) facilitates the use of non-parametric density estimators in conjunction with the implicit copula, which affords plausible approximations to arbitrary bounded margins. In a sequential estimation procedure, the kernel density function is determined first, following the optimisation method proposed by Shimazaki and Shinomoto (2010), which minimises the expected squared error loss of the fitted density.

A considerable body of research has been carried out in an attempt to develop models that perform robustly in the post-global financial crisis period. Various shadow rate models have been proposed (Kim and Singleton, 2012; Krippner, 2013; Bauer and Rudebusch, 2016a; Priebsch et al., 2017), such as the shadow-rate term structure model (Wu and Xia, 2016) and the shadow-rate arbitrage-free Nelson–Siegel model (Christensen and Rudebusch, 2014, 2016). With much of the current literature primarily concerned with the implications for monetary policy, there is a relatively small body of research that pays particular attention to the applications in a forecasting context. Distinguished from the shadow-rate concept, the non-Gaussian macro-yields model proposed in this paper provides an efficient means of handling yields' aggregating around low levels without pre-assuming the existence and location of a possible effective lower bound, of which no clear consensus has been achieved in literature or practice to date. This copula-based approach naturally produces more realistic and reasonable forecasts that concentrate around historical ranges, as well as plausibly adapting to the changing interest rate environment.

4.3 Model Specification

Dynamic term structure models can be regarded as special cases of the DFM, linking the observed yields to the underlying factors by applying either asset pricing theory or empirical numerical techniques. This section begins by providing a brief review of the dynamic term structure modelling, and then proceeds to consider the construction and estimation of copulas in order to arrive at the non-Gaussian macro-yields model.

4.3.1 Macro-Finance Term Structure Modelling

An important strand of the literature takes into account the data-rich environment and attempts to extract pertinent information from large macroeconomic databases (Stock and Watson, 2002; Bernanke and Boivin, 2003; Giannone et al., 2004; Bernanke et al., 2005; Favero et al., 2005; Mönch, 2008; Favero et al., 2012; Exterkate et al., 2013). As presented in Chapter 3, the dynamic factor modelling approach used in conjunction with the Nelson–Siegel representation (Coroneo et al., 2016) explicitly addresses the cross-sectional and serial relationships among a relatively large collection of macroeconomic and yield variables. Revisiting Equations (3.6–3.8), this macro-yields model characterises the dynamics of the cross section of bond yields $Y_{\tau,t}$ and a large panel of macroeconomic variables $Y_{M,t}$ by a condensed state space, consisting of yield curve factors $F_{Y,t}$ and macro-related factors $F_{M,t}$

$$\begin{pmatrix} Y_{\tau,t} \\ Y_{M,t} \end{pmatrix} = \begin{pmatrix} c_Y \\ c_M \end{pmatrix} + \begin{pmatrix} \Gamma_{YY} & \Gamma_{YM} \\ \Gamma_{MY} & \Gamma_{MM} \end{pmatrix} \begin{pmatrix} F_{Y,t} \\ F_{M,t} \end{pmatrix} + e_t, \tag{4.1}$$

with the underlying states F_t following a first-order VAR process

$$\begin{pmatrix} F_{Y,t} \\ F_{M,t} \end{pmatrix} = \begin{pmatrix} \mu_Y \\ \mu_M \end{pmatrix} + \begin{pmatrix} \Pi_{YY} & \Pi_{YM} \\ \Pi_{MY} & \Pi_{MM} \end{pmatrix} \begin{pmatrix} F_{Y,t-1} \\ F_{M,t-1} \end{pmatrix} + \eta_t, \tag{4.2}$$

where $\eta_t \sim N(0, \Sigma_\eta)$, and the idiosyncratic components e_t following an independent autoregressive process

$$e_t = \Pi_e e_{t-1} + \xi_t,\ \xi_t \sim N(0, \Sigma_\xi). \tag{4.3}$$

The yield factors $F_{Y,t}$ and their corresponding factor loadings Γ_{YY} are assumed to follow the parsimonious Nelson–Siegel approach that validates the level, slope, and

curvature interpretations

$$F_{Y,t} = \{F_{\text{Level}}, F_{\text{Slope}}, F_{\text{Curvature}}\}_t, \; c_Y = 0, \tag{4.4}$$

$$\Gamma_{YY} = \Gamma_{NS} = \begin{pmatrix} 1 & \frac{1-e^{-\lambda\tau_1}}{\lambda\tau_1} & \frac{1-e^{-\lambda\tau_1}}{\lambda\tau_1} - e^{-\lambda\tau_1} \\ 1 & \frac{1-e^{-\lambda\tau_2}}{\lambda\tau_2} & \frac{1-e^{-\lambda\tau_2}}{\lambda\tau_2} - e^{-\lambda\tau_2} \\ & \cdots & \cdots \\ 1 & \frac{1-e^{-\lambda\tau_n}}{\lambda\tau_n} & \frac{1-e^{-\lambda\tau_n}}{\lambda\tau_n} - e^{-\lambda\tau_n} \end{pmatrix}. \tag{4.5}$$

The unspanned macro-finance term structure models build directly upon Equation (4.1) to account for the unspanning property, which refers to the notion that part of the macroeconomic information bears no direct relation to the contemporaneous yield curve but provides insights into its future movement. The unspanned models impose the unspanning restrictions by placing constraints on the observation equation $\Gamma_{YM} = 0$ to disentangle the cross section of current yields $Y_{\tau,t}$ from the unspanned factors $F_{M,t}$. Besides, the unspanned models retain the conventional structure between observed yields and yield factors through factor loadings Γ_{YY}, and identify causality in macro-finance linkages through the state transition matrix Π_{YM}. The differences between unspanned affine arbitrage-free and dynamic Nelson–Siegel models lie primarily in the specification of the factor loading block Γ_{YY} that associates the observed yields with the latent yield factors.

4.3.2 Construction of the Copula

Copulas decompose any given multivariate distribution into individual margins and dependence structures, which can be specified and estimated sequentially, allowing enhanced flexibility and alleviating the computational burden. The existence of such a copula function C is guaranteed by Sklar's theorem (Sklar, 1959)

$$F(y_1, y_2, \ldots, y_n) = C(u_1 = F_1(y_1), \ldots, u_n = F_n(y_n)), \tag{4.6}$$

where F denotes the joint distribution function of an n-dimensional random vector $Y = (y_1, y_2, \ldots, y_n)$ with margins $F_1, F_2, \ldots, F_n$. The dependency structure among variables, following removal of the marginal effects, is captured entirely by the copula C. The converse statement of Sklar's theorem allows the combination of arbitrary marginal univariate distributions in a manner specified by the copulas in order to yield valid multivariate characterisations. In contrast to explicit copulas that admit closed-form formulas in terms of $u_1, ..., u_n$ in Equation (4.6), for instance the Archimedean copula family, the implicit copulas adopt a different perspective by retrieving implied dependency from existing multivariate distributions, or even

multivariate time series models, which do not necessarily possess analytical closed-form expressions.

Copula functions characterise the dependency between variables on a quantile scale, meaning that performing any strictly increasing transformation on the margins would not change the underlying copula (McNeil et al., 2005). Accordingly, variables Z with given marginal distributions F_Z can be transformed into other continuously distributed variables $Y = F_Y^{-1}(F_Z(Z)) \sim F_Y$, with the underlying joint dependency remaining unchanged. This can be obtained through the probability integral transform that converts given variables into uniformly distributed values, $F_Z(Z) \sim U(0,1)$, and quantile transformation that returns the threshold value y corresponding to a given probability p, $F_Y(y)$: $Pr(Y \leq y) = p$.

Being capable of extracting the latent dependency structure from any multivariate characterisation, the implicit copula approach enables the availability of a wider range of copulas, not restricted to explicit functional forms or existing parametric copulas. As illustrated by Nelsen (2007), in a bivariate case featuring a joint distribution F with individual margins F_x and F_y, where $C(u,v) = F(F_x^{-1}(u), F_y^{-1}(v))$, new bivariate distributions with arbitrary margins G_x and G_y can be constructed using $C(G_x(x), G_y(y))$ according to the converse statement of Sklar's theorem, with the new distributions inheriting the dependence structure implied by F.

In the context of yield curve modelling and forecasting, the macro-yields model outlined by Coroneo, Giannone, and Modugno (2016) has demonstrated satisfactory performance prior to the global financial crisis by adopting a dynamic factor framework that effectively captures both serial and cross-sectional dependence in a parsimonious manner. Their model, however, implies normally distributed margins and inevitably produces symmetric forecast densities, which seem to be at odds with the empirical evidence of non-Gaussianity in yields, especially in the post-recession period. Thus, the model proposed here aims to borrow only the dependence structure that underlies the macro-yields model, while fitting non-parametric non-Gaussian margins. This can be achieved by extracting the implicit copula that inherits only the dependency of the macro-yields model, which has the potential to generate more accurate forecasts when combined with appropriate choices of marginal distributions.

The macro-yields model (Coroneo, Giannone, and Modugno, 2016) summarises the co-movement of a large collection of yields and macroeconomic variables into a few autoregressive factors F_t while accommodating some degree of autocor-

relation in the idiosyncratic disturbances e_t

$$Z_t = c + \Gamma F_t + e_t, \tag{4.7}$$

$$F_t = \mu + \Pi F_{t-1} + \eta_t,\ \eta_t \sim N(0, \Sigma_\eta), \tag{4.8}$$

$$e_t = \Pi_e e_{t-1} + \xi_t,\ \xi_t \sim N(0, \Sigma_\xi), \tag{4.9}$$

with Z_t denoting the base process specified by Equations (4.7–4.9), that is the multivariate Gaussian state-space time series. Augmenting F_t with e_t casts the above model into the standard state-space formulation

$$Z_t = CX_t + \varepsilon_t, \varepsilon_t \sim N(0, \Sigma_\varepsilon), \tag{4.10}$$

$$X_t = AX_{t-1} + \upsilon_t, \upsilon_t \sim N(0, \Sigma_\upsilon), \tag{4.11}$$

which implies normality for the marginal distributions of Z_t.

The cross-sectional and serial dependence underlying the d-dimensional multivariate time series $Z_t = (Z_{1,t}, \ldots, Z_{d,t})'$ is characterised by the copula density

$$c(u) = \frac{f_Z(z)}{\prod_{t=1}^{T} \prod_{i=1}^{d} f_i(z_{i,t})}, \tag{4.12}$$

which captures the dependence structure in the overall joint distribution $f_Z(z)$ after removing the marginal effects $f_i(z_{i,t})$. Essentially, the implicit copula constructed through the inversion of a multivariate Gaussian state-space representation is the Gaussian copula, and thus is parameterised in terms of the correlation matrix. Consider more generally a stationary autoregressive process of order p for the latent dynamic factors X_t in Equation (4.11), the autocovariance structure of the base process Z_t indicated by Equation (4.10) is

$$\Sigma_Z = \begin{pmatrix} \Sigma_Z(0) & \Sigma_Z(1) & \ldots & \Sigma_Z(p) \\ \Sigma_Z'(1) & \Sigma_Z(0) & \ldots & \Sigma_Z(p-1) \\ \vdots & \ddots & \ddots & \vdots \\ \Sigma_Z'(p) & \Sigma_Z'(p-1) & \ldots & \Sigma_Z(0) \end{pmatrix}, \tag{4.13}$$

where $\Sigma_Z(\tau) = \mathrm{E}[(Z_t - \mathrm{E}[Z_t])(Z_{t-\tau} - \mathrm{E}[Z_{t-\tau}])'] = C\Sigma_X(\tau)C' + \Sigma_\varepsilon, \tau = 0, \ldots, p$. The covariance matrix $\Sigma_X(\tau)$ can be obtained from $\sum_{m=0}^{\infty} \tilde{A}^{m+\tau} \Sigma_{\tilde{v}} (\tilde{A}^m)'$, with $\tilde{A}$ and $\Sigma_{\tilde{v}}$ denoting the autoregressive coefficients and innovation covariance of the equivalent VAR(1) form. The implied dependency of Z_t can be represented by the set of pairwise correlations $\rho_Z(i, j, \tau)$ that characterise the bivariate normal distributions

of all pairs of variables $(Z_{i,t}, Z_{j,t-\tau})$

$$\Sigma_Z = \{\rho_Z(i,j,\tau), i,j = 1,\ldots,d; \tau = 0,\ldots,p\}. \tag{4.14}$$

$\rho_Z(i,j,\tau)$ corresponds to the (i,j)th element of the lag-τ block $\Sigma_Z(\tau)$ in Equation (4.13), addressing the serial dependence $\{\rho_Z(i,j,\tau), \tau > 0\}$ as well as capturing the cross-sectional dependence $\{\rho_Z(i,j,\tau), i \neq j\}$.

As discussed previously, transforming $Z_t = (Z_{1,t},\ldots,Z_{d,t})'$ by means of the inverse cumulative distribution function, that is $Y_{i,t} = F_{Y_i}^{-1}(\Phi_i(Z_{i,t}))$, forms a new multivariate time series $Y_t = (Y_{1,t},\ldots,Y_{d,t})'$ that features certain marginal distributions F_{Y_i} while inheriting the dependency properties from Z_t. More specifically, the pairwise correlations in Y_t are given by

$$\rho_Y(i,j,\tau) = \frac{\mathrm{E}[Y_{i,t}Y_{j,t-\tau}] - \mathrm{E}[Y_{i,t}]\mathrm{E}[Y_{j,t-\tau}]}{\sqrt{\mathrm{Var}[Y_{i,t}]\mathrm{Var}[Y_{j,t-\tau}]}} = \frac{\mathrm{E}[Y_{i,t}Y_{j,t-\tau}] - \mu_i^y\mu_j^y}{\sigma_i^y\sigma_j^y}, \tag{4.15}$$

$$= \frac{\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} F_{Y_i}^{-1}(\Phi_i(z_{i,t}))F_{Y_j}^{-1}(\Phi_j(z_{j,t-\tau}))f_{\rho_z(i,j,\tau)}(z_{i,t},z_{j,t-\tau})dz_{i,t}dz_{j,t-\tau} - \mu_i^y\mu_j^y}{\sigma_i^y\sigma_j^y}, \tag{4.16}$$

where

$$f_{\rho_z(i,j,\tau)}(z_{i,t},z_{j,t-\tau}) = \frac{1}{2\pi\sigma_i^z\sigma_j^z\sqrt{1-\rho_z^2}}\exp\left(-\frac{\frac{(z_{i,t}-\mu_i^z)^2}{(\sigma_i^z)^2} - \frac{2\rho_z(z_{i,t}-\mu_i^z)(z_{j,t-\tau}-\mu_j^z)}{\sigma_i^z\sigma_j^z} + \frac{(z_{j,t-\tau}-\mu_j^z)^2}{(\sigma_j^z)^2}}{2(1-\rho_z^2)}\right), \tag{4.17}$$

with $\mu_{i,j}^{y,z}$ and $\sigma_{i,j}^{y,z}$ denoting their respective means and variances. As suggested by Equation (4.16), the autocorrelation structure of Y_t implied in $\rho_Y(i,j,\tau)$ is determined by $\rho_Z(i,j,\tau)$ after removal of the marginal properties, though analytical forms for $\rho_Y(i,j,\tau)$ are not always readily available (Biller and Nelson, 2003; Biller, 2009).

To summarise, the existence of the copula function is generally guaranteed by the Sklar's Theorem

$$F(z_1,\ldots,z_n;\theta_C) \to C(\Phi_1(z_1),\cdots,\Phi_n(z_n);\theta_C), \tag{4.18}$$

while in this specific case the implicit copula can be expressed in terms of a set of pairwise correlations $\theta_C = \{\rho_Z(i,j,\tau)\}$, as demonstrated in Equations (4.12–4.14). According to the converse statement of Sklar's theorem, flexible marginal density

estimators F_{Y_i} can be employed in conjunction with this implied copula

$$F(y_1, \ldots, y_n; \theta_C) \leftarrow C(F_{Y_1}(y_1), \cdots, F_{Y_n}(y_n); \theta_C) \tag{4.19}$$

to construct a joint distribution that accommodates the non-Gaussian characteristics of univariate series while addressing the complex relationships between variables, as described by Equations (4.15–4.17).

Sharing the same implicit copula, the macro-yields model described in Equations (4.1–4.3) is structurally similar to the multivariate representation defined in Equations (4.7–4.19), but a crucial difference lies in whether the non-Gaussianity of variables is treated. While the former implies normal distributions Φ for all latent variables Z_t, regardless of their physical properties, the latter relaxes this assumption to permit any form of distributions F for any subset of observable variables Y_t, including normal distributions with which the original model is reproduced. In doing so, this copula-based approach encompasses the macro-yields model as a special case, but accommodates a broad range of model variations that improve the forecasting accuracy through enhanced assessment of marginal behaviours[1]. To underscore the differences as well as similarities, this copula-based approach is labelled the "non-Gaussian macro-yields" model in the following analysis.

Exploiting the implicit copula is then equivalent to implementing a state-space model. In the E-step of the expectation maximum algorithm, applying filtering and smoothing techniques generates closed-form formulas for the conditional expectation of the conditional log-likelihood. The M-step maximises the expected value using the usual first-order optimality conditions that readily yield analytical formulations. Various model specifications can be achieved through imposing linear constraints in a form analogous to $C_1 \text{vec}\tilde{\Gamma} = h_1$ and $C_2 \text{vec}\tilde{\Pi} = h_2$, in which case closed-form solutions are still available, as exemplified by Wu, Pai, and Hosking (1996). Both filtered states and optimised parameters are updated iteratively until convergence, rendering consistent estimates of the state-space model and thus the implicit copula.

The canonical maximum likelihood method enables the use of non-parametric estimation of the cumulative density function in conjunction with the implicit copula, providing a convenient and flexible means of dealing with the non-Gaussian features in margins (see Appendix 4.A). The kernel density estimator is a non-

[1]This generalisation could introduce a few differences in other respects. The original model imposes the Nelson–Siegel structure on the observation equation and specifies zero blocks in the factor loading matrix to account for the unspanning properties. In the implementation of the inversion copula, however, the Nelson–Siegel parameters do not necessarily provide the optimum description of links between latent variables Z_t and their corresponding common factors F_t. Unspanning properties can still be specified by constraining Γ_{YM} to be zero in the underlying state-space model.

parametric density estimator that aggregates the component smoothing function k at each datum Y_i

$$\hat{f}_w(y) = \frac{1}{nw} \sum_{i=1}^{n} k(\frac{y - Y_i}{w}), \tag{4.20}$$

producing reasonably accurate characterisations of various distributions. Different choices of kernel functions k exhibit different efficiencies, but do not produce substantially different estimates, among which the Gaussian kernel

$$k(u) = \frac{1}{\sqrt{2\pi}} \exp(-\frac{u^2}{2}) \tag{4.21}$$

has proven reliable in providing plausible approximations to arbitrary bounded densities. Once the kernel function is determined, only the kernel bandwidth w remains as a free parameter, which is responsible for the goodness-of-fit of the density estimate to the observed data. The bandwidth parameter controls the smoothness of the fitted curve and has profound consequences for the model performance. This study follows the work of Shimazaki and Shinomoto (2010), which selects the optimal bandwidth by reducing the expected squared loss of the overall density fitting. In out-of-sample analysis, kernel densities are fitted recursively for each forecast origin utilising the observations up to time t with the optimal bandwidths selected over the same sample (Appendix 4.B).

4.4 Empirical Analysis

This section demonstrates the application of the non-Gaussian macro-yields model proposed in Section 4.3.2 using Fama-Bliss yields covering maturities from 3 months to 5 years. The first set of analyses in this section establishes the statistical properties of yields and changes in their distributions over time. What follows is a direct assessment of the in-sample and out-of-sample forecasting performance of the non-Gaussian macro-yields model over the entire 1970:M1–2016:M12 sample period that extends to the post-global financial crisis era. The non-Gaussian scheme consistently outperforms the baseline method, that is its Gaussian counterpart, the inferior performance of which is attributable largely to its non-robustness in the post-global financial crisis era.

4.4.1 Marginal Distributions of Yields

Before proceeding to apply forecast evaluation techniques, it would be pertinent to investigate whether the non-Gaussian nature of yields emerges from the particular low-yield environment following the crisis, or whether it has been more prevalent in practice than formerly assumed in the literature for a longer period of time. Normality can be checked by examining higher-order moments such as skewness and kurtosis that describe certain aspects of the shapes of distributions.

To assess how asymmetric and leptokurtic behaviour changes over time, Figure 4.1 reports skewness and kurtosis over expanding windows that employ all available observations from 1970:M1 up to time t. After the recession in the early 1980s, all yields exhibit significant positive skewness (> 0) and excess kurtosis (> 3), which together indicate asymmetric tails extending towards more positive values. These properties can be observed across maturities and are more pronounced for the very front-end of the curve. Tracking variations over time, the higher-order measures climb until a local peak is reached, usually at the onset of recession, and then come down in the subsequent recovery, revealing how yield distributions evolve over the business cycles.

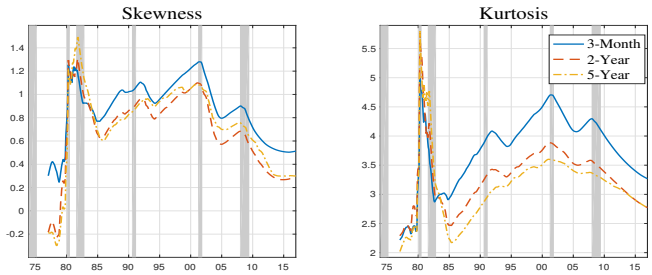

Figure 4.1: Time-varying skewness and kurtosis. This figure shows the skewness and kurtosis calculated recursively using observations from 1970:M1 to current time t along the x-axis. 3-month (solid blue line), 2-year (dashed red line), and 5-year (dash-dotted yellow line) yields are displayed to represent the behaviour of the short-, medium-, and long-end of the curve. The evolution of business cycles is indicated by the grey-shaded areas that correspond to economic contractions identified by NBER.

For a more rigorous examination of deviations from Gaussianity, the Shapiro-Wilk (Shapiro and Wilk, 1965) or Shapiro-Francia (Shapiro and Francia, 1972) tests for normality, adjusted by Royston (1992, 1993, 1995) to handle sample sizes up to 5,000, can be performed on yield series (Appendix 4.C). P-values less than 0.05 can be consistently observed across all maturities over time, implying departures from the normal distribution. Figure 4.2 displays the marginal densities fitted with the kernel estimator described earlier (solid blue line). This approach gives the non-parametrically estimated marginal distributions for the subsequent analysis with copulas models. Comparison of the marginal probability density estimates gener-

ated by kernel distributions, which conform well to the data, and those by normal distributions (dashed red line) with equal mean and variance, further points towards the presence of non-Gaussian features.

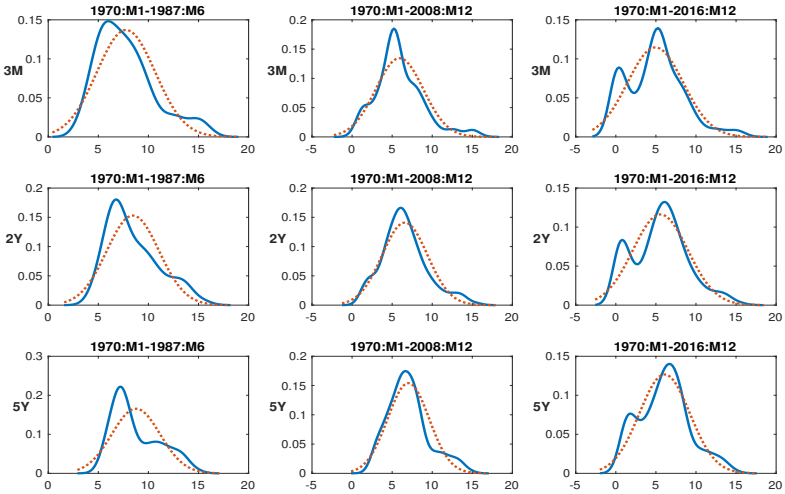

Figure 4.2: Marginal densities fit to yields. This figure compares the marginal densities fit to 3-month (upper), 2-year (middle), and 5-year (lower) yields using nonparametric kernel density functions (solid blue line) and normal probability density functions (dashed red line) over expanding time periods: 1970:M1–1987:M6 (left), –2008:M12 (middle), and –2016:M12 (right).

Corroborating earlier findings, kernel density curves exhibit apparent asymmetry and leptokurtosis in the yields of all maturities across different time frames, in particular when extending the sample period from 1970:M1–2008:M12 to 1970:M1–2016:M12, whereby the inclusion of the post-global financial crisis observations, which aggregate around low-yield values, transforms the unimodal distributions (middle panel) to bimodal ones (right panel). Overall, normality has been vigorously challenged by statistical measures and normality tests, particularly at the very front-end of the curve and during periods of recession, cementing our concerns regarding the prevailing Gaussian assumptions.

4.4.2 In-Sample and Out-of-Sample Predictability

To assess the ability of the copula-based approach to forecast in the presence of non-Gaussian characters, this section presents a comparison of yield forecasts produced by the non-Gaussian schemes and their Gaussian counterparts proposed by Coroneo, Giannone, and Modugno (2016). Tables 4.1–4.6 detail a practical implementation of models utilising the original macroeconomic dataset that consists of 14 variables measuring various aspects of aggregate economic activity (see Appendix 3.A), to examine the in-sample and out-of-sample predictability regarding the yield curve and excess bond returns.

Table 4.1: Out-of-sample predictability of the yield curve (2016:M12)

	MY						non-Gaussian MY					
Maturity	3m	1y	2y	3y	4y	5y	3m	1y	2y	3y	4y	5y
h=1	0.047	0.057	0.073	0.082	0.085	0.083	**0.044**	**0.050**	**0.071**	**0.078**	**0.084**	**0.080**
h=3	0.250	0.299	0.336	0.335	0.326	0.309	**0.175**	**0.209***	**0.275***	**0.278****	**0.284***	**0.271***
h=6	0.797	0.831	0.825	0.770	0.724	0.680	**0.432***	**0.496***	**0.559****	**0.537****	**0.538****	**0.516****
h=12	2.491	2.332	2.057	1.790	1.582	1.402	**1.090****	**1.182****	**1.124****	**0.998****	**0.919****	**0.855****
h=24	5.698	5.284	4.623	4.028	3.501	3.102	**3.059***	**3.294***	**2.854***	**2.394***	**2.067***	**1.836***
	OY						non-Gaussian OY					
Maturity	3m	1y	2y	3y	4y	5y	3m	1y	2y	3y	4y	5y
h=1	0.041	0.055	0.075	0.082	0.085	0.083	**0.036***	**0.050**	**0.068***	**0.076****	**0.081****	**0.078****
h=3	0.200	0.271	0.331	0.326	0.306	0.291	**0.163***	**0.238**	**0.291****	**0.292****	**0.288***	**0.275***
h=6	0.614	0.732	0.774	0.711	0.638	0.600	**0.495****	**0.640***	**0.664****	**0.621****	**0.592**	**0.568**
h=12	2.026	2.091	1.938	1.659	1.401	1.246	**1.628****	**1.834****	**1.656****	**1.427****	**1.258***	**1.156**
h=24	5.582	5.385	4.770	4.042	3.400	2.910	**4.540****	**4.870****	**4.265****	**3.597****	**3.119****	**2.738**

Note. This table reports the MSFE of the Gaussian (left) and non-Gaussian (right) versions of macro-yields model (upper) and only-yields model (lower) for 3m-, 1y-, 2y-, 3y-, 4y-, and 5y-yields (in columns) over horizons of h=1, 3, 6, 12, and 24 months (in rows). Using a recursive scheme, forecasts are constructed using monthly data from 1970:M1 and are evaluated over 1990:M1–2016:M12 as the observation window expands. The bold numbers indicate the cases where the non-Gaussian approach outperforms its Gaussian counterpart in terms of out-of-sample forecasts of the yield curve. * and ** mark significant outperformance of the non-Gaussian approach with respect to its Gaussian rival at 10% and 5% level, according to the standard reality check with 10000 bootstrap replications and a circular block bootstrap with window size of 12 (see Appendix 4.D).

Out-of-sample forecasts over 1–24-month horizons are constructed by adopting a recursive forecasting scheme that employs monthly data available from 1970:M1 up to the evaluation point, which spans from 1990:M1. As the observation window expands alongside the sample size, both the margins and the copulas are re-estimated sequentially until the end of the evaluation period 2016:M12. Over the full period, practices that account for non-Gaussianity are found to deliver consistently superior predictions at all horizons across the full spectrum compared to those restricted to the Gaussian case, as can be seen from Table 4.1 that reports the mean squared forecast error (MSFE).

Examining the performance with respect to the random walk (Table 4.2) finds that inducing non-Gaussianity by adopting copula techniques leads to a reduction of 0.13 in rMSFE for the only-yields model and an even larger reduction of 0.39 for the macro-yields model. The most notable decreases occur at the short-end of the yield curve for multi-step-ahead forecasting; for example, the non-Gaussian macro-yields model reduces the 1-year-ahead rMSFE for the 3-month yield to less than half of the error that would have resulted from the plain macro-yields model, from 1.345 to 0.588.

Table 4.2: Out-of-sample predictability of the yield curve (rMSFE)

	MY						non-Gaussian MY					
Maturity	3m	1y	2y	3y	4y	5y	3m	1y	2y	3y	4y	5y
h=1	1.085	1.110	1.116	1.103	1.066	1.058	1.026	0.967	1.078	1.043	1.046	1.020
h=3	1.231	1.240	1.231	1.189	1.169	1.156	0.862	0.865*	1.006	0.985	1.020	1.012
h=6	1.347	1.345	1.365	1.338	1.311	1.288	0.731**	0.804**	0.925	0.934	0.975	0.977
h=12	1.345	1.328	1.394	1.433	1.443	1.406	0.588**	0.673**	0.762**	0.800**	0.838*	0.857*
h=24	1.192	1.192	1.353	1.511	1.616	1.735	0.640**	0.743**	0.835*	0.898	0.954	1.027
	OY						non-Gaussian OY					
Maturity	3m	1y	2y	3y	4y	5y	3m	1y	2y	3y	4y	5y
h=1	0.939	1.058	1.149	1.106	1.057	1.053	0.833**	0.960	1.041	1.021	1.008	0.990
h=3	0.983	1.123	1.212	1.158	1.099	1.090	0.802*	0.987	1.063	1.035	1.033	1.028
h=6	1.038	1.185	1.280	1.237	1.156	1.137	0.836*	1.037	1.099	1.081	1.073	1.077
h=12	1.094	1.190	1.313	1.329	1.278	1.250	0.879	1.044	1.122	1.143	1.148	1.159
h=24	1.168	1.214	1.396	1.516	1.569	1.628	0.950	1.098	1.248	1.349	1.440	1.532

Note. This table reports the rMSFE of the Gaussian (left) and non-Gaussian (right) versions of macro-yields model (upper) and only-yields model (lower) with respect to the random walk for 3m-, 1y-, 2y-, 3y-, 4y-, and 5y-yields (in columns) over horizons of h=1, 3, 6, 12, and 24 months (in rows). Using a recursive scheme, forecasts are constructed using monthly data from 1970:M1 and are evaluated over 1990:M1–2016:M12 as the observation window expands. * and ** mark significant outperformance with respect to the random walk at 10% and 5% level, according to the standard reality check with 10000 bootstrap replications and a circular block bootstrap with window size of 12.

One striking observation to emerge from Table 4.2 is that over this extended period, the incorporation of macroeconomic variables restricts the Gaussian model's ability to forecast rather than fostering it, as illustrated by the case in the 1-year-ahead forecast for the 1-year yield (in row 4 and column 2 of each panel), where the rMSFE increases from 1.190 (OY) to 1.328 (MY). Facilitated by the use of copulas that take into account non-Gaussianity, the macroeconomic components come into play and enable remarkable improvements in yield forecasting, as indicated by the reduction in rMSFE from 1.044 (non-Gaussian OY) to 0.673 (non-Gaussian MY). Comparing across models in Table 4.2, combining the benefits of copula-based methods and macroeconomic information (i.e. the non-Gaussian macro-yields model) delivers forecasts with superior accuracy and demonstrates improved, or at least comparable, performance relative to the random walk across maturities at virtually all but one-step forward.

Before turning to the out-of-sample predictability of excess bond returns, Table 4.3 examines the in-sample R^2 for 1-year-ahead 1-year-holding-period excess returns on bonds with maturities ranging from 2 to 5 years. With regard to excess bond returns, powerful predictors of perceived usefulness and popularity include the CP factor (Cochrane and Piazzesi, 2005) constituted by forward rates, and the LN factor (Ludvigson and Ng, 2009) extracted from 131 monthly macroeconomic indi-

cators. These factors contribute to predictability by providing valuable information regarding the sovereign debt market and overall macroeconomic conditions. For ease of reproducibility, this study re-constructs the LN factor[2] using the FRED-MD (McCracken and Ng, 2016) database by iteratively conducting PCA on observations available up to time t. The predictability is then assessed according to the two-step procedures outlined in Cochrane and Piazzesi (2005) and Ludvigson and Ng (2009).

Owing to the adoption of the non-Gaussian setup, significant increases in R^2 are recorded across all maturities, producing enhanced in-sample performance (marked in bold in Table 4.3). Closer inspection of the comparison between the "pure-yields" approaches (i.e. OY, non-Gaussian OY, and CP) and their "macro-finance" counterparts (i.e. MY, non-Gaussian MY, and LN+CP) reveals the importance of incorporating macroeconomic information for predicting bond returns. Careful treatment of non-Gaussianity results in modest but consistent improvement in terms of in-sample performance when compared with the Gaussian macro-yields model, in particular for the 5-year bond. Taking a slightly different perspective that compares term structure models with predictive regressions, the non-Gaussian macro-yields model outperforms LN+CP by 8.2% on average in terms of R^2, most notably in the 2-year bond (16.20%), which suggests a role for the term structure representation in forecasting bond returns.

Table 4.3: In-sample predictability of excess bond returns (2016:M12)

Maturity	OY	non-Gaussian OY	MY	non-Gaussian MY	CP	LN	LN+CP
2y	12.35%	15.58%	40.36%	**45.82%**	15.58%	22.26%	29.62%
3y	13.21%	16.00%	38.35%	**40.85%**	18.09%	24.78%	33.52%
4y	14.70%	19.13%	34.66%	**39.55%**	21.05%	25.22%	36.07%
5y	15.89%	21.51%	31.37%	**39.89%**	18.69%	25.10%	34.22%

Maturity	MY - OY (Gaussian)	MY - OY (non-Gaussian)	LN+CP - CP
2y	28.00%	30.23%	14.04%
3y	25.14%	24.85%	15.43%
4y	19.96%	20.42%	15.02%
5y	15.48%	18.37%	15.53%

Note. This table reports the in-sample R^2 for 1-year-ahead forecasts of 1-year-holding-period excess returns on 2- to 5-year bonds (in rows) in the full 1970:M1–2016:M12 sample. Models examined include Gaussian state-space models: the only-yields model (OY), macro-yields model (MY); their non-Gaussian counterparts: the non-Gaussian OY, non-Gaussian MY; and the predictive regressions on prevailing factors: the CP factor, LN factor and their combination LN+CP. The lower panel organises the results by subtracting R^2s of the pure-yields models from those of their macro-finance counterparts. The bold numbers mark the best in-sample performance in terms of R^2.

Moving on to consider the out-of-sample predictability, Table 4.5 presents

[2]Though being updated by Ludvigson and Ng (2009), their LN factor is extracted by performing PCA on their whole sample 1960M3:2018M6, making it unsuitable for out-of-sample applications.

comparative results in terms of rMSFE with respect to the expectations hypothesis, with the non-Gaussian schemes outperforming their Gaussian counterparts in all instances (also see Table 4.4). Taking into account the possible non-Gaussianity results in an average reduction of 0.15 in rMSFE produced by the only-yields model and a considerable reduction of 0.59 for the macro-yields model. Generating the lowest forecast errors (marked in bold), the non-Gaussian macro-yields model possesses competitive advantages over its rivals examined in Table 4.5, even when compared with LN+CP that combines perspectives provided by forward rates and large macroeconomic panels. The superiority becomes more pronounced in the case of the 2-year bond (0.96 in LN+CP vs 0.70 in non-Gaussian MY), and is consistently observed across maturities with a marked decrease of 0.16 in rMSFE on average.

Table 4.4: Out-of-sample predictability of excess bond returns (2016:M12)

	OY	non-Gaussian OY	MY	non-Gaussian MY
2y	2.09	**1.83****	2.33	**1.18****
3y	7.75	**6.62****	8.23	**4.50****
4y	14.93	**12.84****	16.11	**8.99****
5y	22.41	**20.13***	25.31	**14.70****

Note. This table reports the MSFE for 1-year-ahead forecasts of 1-year-holding-period excess returns on 2- to 5-year bonds. The bold numbers indicate the cases where the non-Gaussian approach outperforms its Gaussian counterpart, with * and ** marking significant outperformance at 10% and 5% level.

Table 4.5: Out-of-sample predictability of excess bond returns (rMSFE)

Maturity	OY	non-Gaussian OY	MY	non-Gaussian MY	CP	LN	LN+CP
2y	1.23	1.08	1.37	**0.70****	1.32	1.22	0.96
3y	1.23	1.05	1.30	**0.71****	1.40	1.07	0.90
4y	1.18	1.01	1.27	**0.71****	1.41	0.97	0.84
5y	1.13	1.02	1.28	**0.74****	1.40	0.90	0.82

Note. This table reports the rMSFE with respect to the expectations hypothesis. The bold numbers mark the best out-of-sample predictive performance for each maturity in terms of rMSFE. * and ** mark significant outperformance with respect to the expectations hypothesis at 10% and 5% level.

The manner in which the macroeconomic variables are incorporated, that is, in a Gaussian or non-Gaussian setting, plays a critical role in exploiting the informational content of macroeconomic data when modelling and forecasting the yield curve. Table 4.6 re-examines Table 4.5 by placing particular focus on comparison of the "pure-yields" approaches (i.e. OY, non-Gaussian OY, and CP) and their "macro-finance" counterparts (i.e. MY, non-Gaussian MY, and LN+CP) in terms of out-of-sample forecasting. While the Gaussian MY model underperforms the Gaussian OY model by an average of 0.11 across maturities in terms of rMSFE, the non-Gaussian MY enhances the prediction of the non-Gaussian OY by 0.32 on

average. This finding collaborates with that found in Table 4.2 to further support the notion that how macroeconomic components are utilised can make a material difference. More importantly, outperforming the non-predictable expectations hypothesis is achievable only when the macro-finance linkage is well attended to, as per the non-Gaussian MY and CP+LN. Appropriate characterisations of non-Gaussianity further contribute to forecast accuracy, with the non-Gaussian MY producing superior out-of-sample results marked by the lowest rMSFE values.

Table 4.6: Role of macroeconomic variables in yield forecasting

Maturity	OY	vs	MY	non-Gaussian OY	vs	non-Gaussian MY	CP	vs	LN+CP
2y	1.23	$<$	1.37	1.08	$>$	0.70	1.32	$>$	0.96
3y	1.23	$<$	1.30	1.05	$>$	0.71	1.40	$>$	0.90
4y	1.18	$<$	1.27	1.01	$>$	0.71	1.41	$>$	0.84
5y	1.13	$<$	1.28	1.02	$>$	0.74	1.40	$>$	0.82

Note. This table compares the "pure-yields" approaches (i.e. OY, non-Gaussian OY, and CP) and their "macro-finance" counterparts (i.e. MY, non-Gaussian MY, and LN+CP) in terms of rMSFE for 1-year-ahead forecasts of 1-year-holding-period excess returns on bonds with maturities 2–5y.

Overall, exploiting the informational content of the macroeconomic dataset by means of copulas that handle non-Gaussianity in marginal yield distributions exhibits improved predictive ability compared to the vanilla macro-yields model, the expectations hypothesis and other existing approaches. Considering that the non-Gaussian schemes employ the same implicit copula as well as the same information set, with the most important difference lying in the treatment of yield variables, this substantial enhancement of forecasts can be soundly attributed to the new approach's successful exploitation of the non-Gaussian attributes inherent in yield series. This is especially the case for the period after the recent recession. During this period, while the dramatic macroeconomic fluctuations largely propagate to the yield dynamics to drive yields beyond their usual ranges as in the macro-yields model, the new approach plausibly adapts to the changing interest rate environment, as demonstrated in Figure 4.2.

4.4.3 Evaluating Performance over Time

With the non-Gaussian macro-yields model, considerable enhancements have been observed in the average forecast error (measured by MSFE or rMSFE) over the out-of-sample 1990:M1–2016:M12 evaluation period in Section 4.4.2. This subsection conducts a more detailed investigation of the prediction performance over time via a graphical representation, providing additional insight into patterns of improvement and deterioration as the business cycle evolves.

To provide a visual assessment, Figure 4.3 and Figure 4.4 compare the

squared forecast errors obtained by the non-Gaussian macro-yields model and other existing methods, with the forecast origin travelling along the x-axis. Corroborating earlier findings presented in Section 4.4.2, Figure 4.3 shows that while addressing non-Gaussianity through copulas introduces a modest improvement in the only-yields model (right panel), it contributes substantially to the forecast accuracy of the macro-yields model (left panel) for the majority of the time, most predominantly in the post-global financial crisis era. This enhanced predictability can be observed across the entire maturity spectrum, from 3-month yield (top-left panel) to 5-year yield (bottom-left panel). Compared to the random walk, with regard to 1-year-ahead forecasts, outperformance of the non-Gaussian macro-yields model primarily occurs in the first half of the 1990s and the period between recessions, such as the 2002-2007 period. During economic downturns, the non-Gaussian scheme remains superior for the short-end but is slightly inferior to the random walk at the long-end, which is in alignment with the findings in Table 4.2.

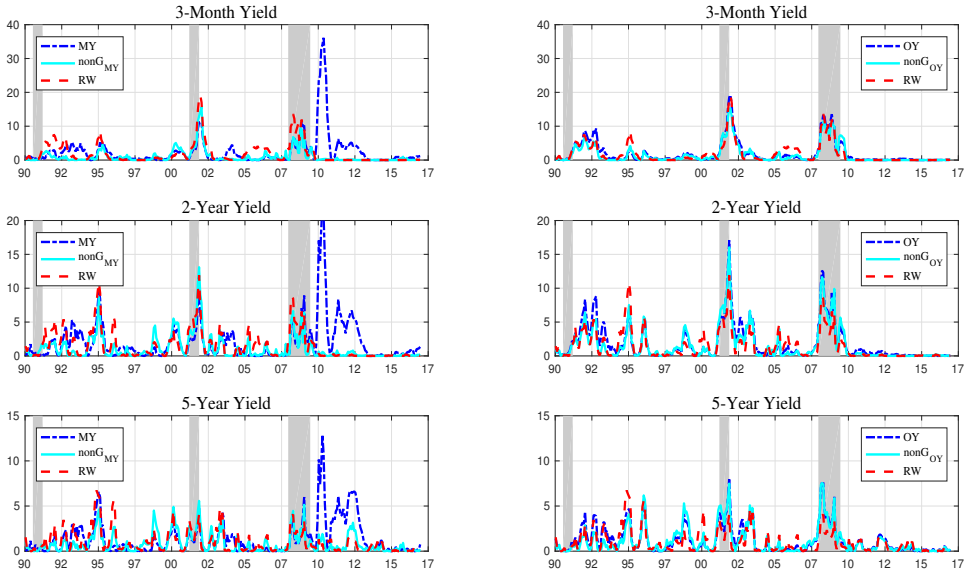

Figure 4.3: Forecasting the yield curve over time. This figure displays the squared errors of 1-year-ahead forecasts for 3m-(upper), 2y-(middle), and 5y-yields (lower), with the evaluation dates (the end dates of expanding observation windows) reported along the x-axis. The dash-dotted lines represent the original Gaussian models (MY, OY) while the continuous lines plot their non-Gaussian counterparts (nonG_{MY}, nonG_{OY}), with the dashed lines reporting the random walk benchmark (RW). The evolution of business cycles is indicated by the grey-shaded areas that correspond to economic contractions identified by NBER.

Similarly, with regard to bond-return forecasting, the non-Gaussian treatment adds predictive power to the macro-yields model for the majority of the evaluation of 27 years (Figure 4.4). Even when compared with the expectations hypothesis, the non-Gaussian macro-yields model is competitive over time, which is consistent with its average rMSFE of approximately 0.7 in the full sample.

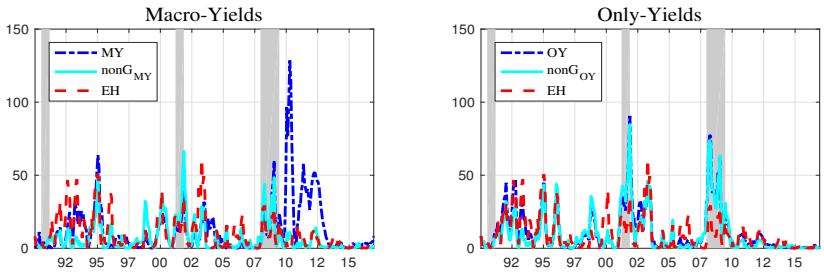

Figure 4.4: Forecasting excess bond returns over time. This figure displays the squared errors of 1-year-ahead forecasts for 1-year-holding-period excess returns, with the evaluation dates (the end dates of expanding observation windows) reported along the x-axis. The dash-dotted lines represent the original Gaussian models (MY, OY) while the continuous lines plot their non-Gaussian counterparts (nonG_{MY}, nonG_{OY}), with the dashed lines reporting the expectations hypothesis benchmark (EH). The evolution of business cycles is indicated by the grey-shaded areas that correspond to economic contractions identified by NBER.

Overall, as suggested by Figure 4.3 and Figure 4.4, the enhancement of predictability owing to the incorporation of non-Gaussianity can be observed over various time intervals. Particularly, the copula-based approach offers a promising solution to handling the low variation observed in yields in the post-global financial crisis period, when low rates persist despite the stabilisation of economic activity and financial conditions. Table 4.7 reports the relative performance of the non-Gaussian model compared to the Gaussian model during the post-global financial crisis period, which produces rather low rMSFE values.

Table 4.7: Non-Gaussian vs Gaussian macro-yields model in the post-GFC period

Maturity	3m	1y	2y	3y	4y	5y
h=1	0.115*	0.134*	0.441**	0.615**	0.806	0.705**
h=3	0.073*	0.078*	0.218**	0.327**	0.552**	0.580**
h=6	0.096*	0.082**	0.163**	0.237**	0.397**	0.445**
h=12	0.073**	0.064**	0.087**	0.123**	0.197**	0.249**
h=24	0.208**	0.213**	0.231**	0.257**	0.289**	0.313**

Note. This table reports the rMSFE of the non-Gaussian macro-yields model with respect to the Gaussian macro-yields model for 3m-, 1y-, 2y-, 3y-, 4y-, and 5y-yields (in columns) over horizons of h=1, 3, 6, 12, and 24 months (in rows) in the post-global financial crisis (GFC) period. * and ** mark significant outperformance of the non-Gaussian approach with respect to its Gaussian rival at 10% and 5% level, according to the standard reality check with 10000 bootstrap replications and a circular block bootstrap with window size of 12.

Avoiding undesirable outcomes that extend far beyond the ordinary bounds, the non-Gaussian scheme sustains advantages over the macro-yields model and delivers plausible predictions not only in a low-yield environment, but also in other circumstances. Revisiting the density estimates employed for the yield margins in Figure 4.5 enables us to understand how the copula techniques achieve improved performance relative to the Gaussian approaches. Figure 4.6 supports this interpre-

tation by examining the predictive densities for yields at different forecast origins when moving from pre- to post-recession periods.

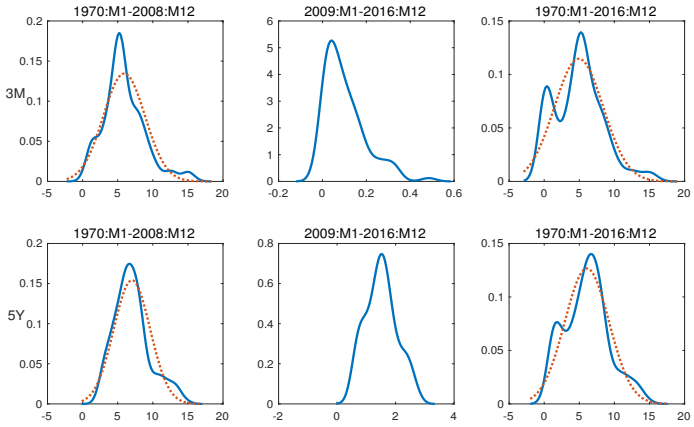

Figure 4.5: Marginal densities fit to yields in 2009–2016. This figure compares the marginal densities fit to 3-month (upper) and 5-year (lower) yields using nonparametric kernel density functions (solid blue line) and normal probability density functions (dashed red line) over time periods: 1970:M1–1987:M6 (left), 2009:M1–2016:M12 (middle), and 1970:M1–2016:M12 (right).

Being capable of addressing the substantial asymmetry, the proposed non-Gaussian model captures the yields' inclination towards more positive values (solid blue line in the upper-right panel of Figure 4.6). In contrast, the Gaussian macro-yields model with symmetric margins assigns higher probabilities to negative forecasts (dashed red line), resulting in severe under-prediction of yields during the crisis period. Extending to the full 1970:M1–2016:M12 sample creates an additional mode in the yield distributions, as exhibited in Figure 4.5, which is well exploited and accounted for by the proposed copula to generate more reasonable forecasts that aggregate around the low yield levels (solid blue line in the lower-left panel of Figure 4.6). The distributions associated with the Gaussian model, however, are constrained to be unimodal and thus imply relatively high-mode densities for post-recession forecasts (dashed red line), leading to over-prediction of yields during the recovery period.

Collectively, these results indicate that the non-Gaussian macro-yields model proposed in this chapter generates plausible forecasting over time, offering a satisfactory solution to the issues raised in Chapter 3, which the linear-Gaussian DFM is incapable of adequately handling. In particular, the non-Gaussian approach overwhelmingly and substantially outperforms the Gaussian setup in the post-global financial crisis period, to which much of its superiority over the entire period can be ascribed.

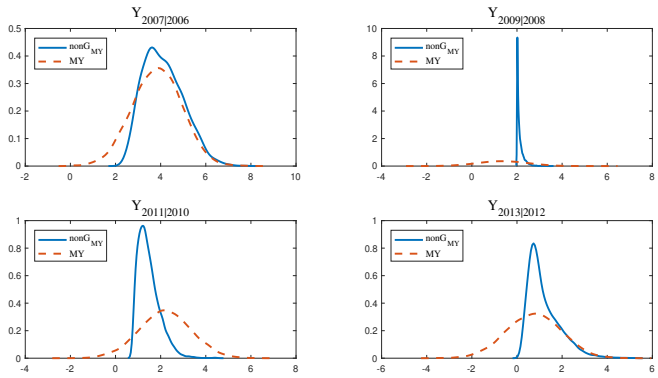

Figure 4.6: Density forecast of yields. This figure displays the 1-year-ahead density forecasts of 5-year yield using non-Gaussian macro-yields model (blue solid line) and macro-yields model (red dashed line) at different forecast origins: 15-Nov-2006 (upper-left), 15-Nov-2008 (upper-right), 15-Nov-2010 (lower-left), 15-Nov-2012 (lower-right). $Y_{t+12|t}$ refers to the 1-year-ahead forecasts at $t+12$ from forecast horizon t.

4.5 Alternative Dataset and Economic Value

4.5.1 Alternative Dataset

Before concluding this chapter, this subsection investigates whether alternative datasets offer opportunities for superior forecasts, especially at one-step-ahead for which the random walk proves to be the most effective. It is, however, beyond the scope of this chapter to provide a comprehensive discussion of big data applications, which requires a detailed investigation and is thus left for future research. This section only represents one of many possibilities in this direction, whereby different datasets can be exploited for differing forecast ranges. In an attempt to utilise the FRED-MD dataset, Table 4.8 revisits previous analyses by replacing the 14 macroeconomic variables with eight common factors extracted from the FRED-MD dataset using PCA, referred to as "FRED-yields", or including an additional cubic term on that basis, denoted as "LN-yields".

In line with previous conclusions, the non-Gaussian schemes significantly improve the forecast accuracy across all maturities and at all horizons. The results with the large dataset compare favourably with those achieved by the original dataset at short-term horizons, reducing the rMSFE by nearly 0.1 on average. The non-Gaussian FRED-yields model produces smaller errors than the random walk for 1- and 3-month-ahead forecasts, while the incorporation of the cubic term (i.e. the non-Gaussian LN-yields model) delivers superior one-step-ahead forecast for the 3-month yield, exemplified by an rMSFE of 0.831.

Table 4.8: Forecasting the yield curve using alternative datasets

	Macro-Yields						non-Gaussian Macro-Yields					
Maturity	3m	1y	2y	3y	4y	5y	3m	1y	2y	3y	4y	5y
h=1	1.085	1.110	1.116	1.103	1.066	1.058	1.026	0.967	1.078	1.043	1.046	1.020
h=3	1.231	1.240	1.231	1.189	1.169	1.156	0.862	0.865	1.006	0.985	1.020	1.012
h=6	1.347	1.345	1.365	1.338	1.311	1.288	0.731	0.804	0.925	0.934	0.975	0.977
h=12	1.345	1.328	1.394	1.433	1.443	1.406	0.588	0.673	0.762	0.800	0.838	0.857
h=24	1.192	1.192	1.353	1.511	1.616	1.735	0.640	0.743	0.835	0.898	0.954	1.027
	FRED-Yields						non-Gaussian FRED-Yields					
Maturity	3m	1y	2y	3y	4y	5y	3m	1y	2y	3y	4y	5y
h=1	1.516	1.286	1.098	0.998	0.894	0.911	**0.972**	**0.968**	**0.954**	**0.939**	**0.937**	**0.922**
h=3	1.601	1.509	1.419	1.290	1.163	1.124	**0.630**	**0.786**	**0.897**	**0.911**	**0.945**	**0.954**
h=6	1.688	1.735	1.725	1.606	1.444	1.374	0.702	0.920	1.019	1.030	1.043	1.055
h=12	1.499	1.533	1.629	1.632	1.554	1.498	0.857	1.043	1.140	1.174	1.195	1.216
h=24	1.344	1.352	1.549	1.694	1.764	1.841	0.995	1.141	1.298	1.428	1.551	1.680
	LN-Yields						non-Gaussian LN-Yields					
Maturity	3m	1y	2y	3y	4y	5y	3m	1y	2y	3y	4y	5y
h=1	1.217	1.128	1.057	1.009	0.940	0.966	**0.831**	**0.911**	**0.977**	**0.979**	**0.973**	**0.952**
h=3	1.339	1.408	1.342	1.259	1.171	1.168	0.721	0.917	1.004	0.999	1.007	1.013
h=6	1.546	1.653	1.662	1.581	1.464	1.441	0.768	0.991	1.082	1.084	1.083	1.095
h=12	1.512	1.569	1.691	1.722	1.679	1.656	0.845	1.026	1.139	1.186	1.210	1.232
h=24	1.340	1.351	1.561	1.726	1.822	1.934	0.927	1.059	1.212	1.337	1.458	1.579

Note. This table reports the rMSFE using different macroeconomic datasets for 3m-, 1y-, 2y-, 3y-, 4y-, and 5y-yields (in columns) over horizons of h=1, 3, 6, 12, and 24 months (in rows) in the full 1970:M1–2016:M12 sample. The 14 macroeconomic variables in the macro-yields (upper-left) model are replaced by the eight principal components extracted from the FRED-MD dataset (middle-left), denoted as FRED-yields, or plus the cube of the first principal component (lower-left), denoted as LN-yields. The corresponding non-Gaussian models are reported accordingly in the right panel. Bold numbers mark outperformance relative to the non-Gaussian macro-yields model (upper-right).

Table 4.9: Forecasting excess bond returns using alternative datasets

In-Sample Predictability of Excess Bond Returns (R^2)						
Maturity	Macro-yields	nonG$_{\text{MY}}$	FRED-yields	nonG$_{\text{FRED-Y}}$	LN-yields	nonG$_{\text{LN-Y}}$
2y	40.36%	45.82%	14.88%	22.37%	25.75%	19.26%
3y	38.35%	40.85%	13.90%	22.61%	24.56%	19.24%
4y	34.66%	39.55%	14.22%	24.72%	23.10%	21.56%
5y	31.37%	39.89%	14.88%	27.15%	23.65%	24.60%

Note. This table reports the in-sample R^2 for the 1-year-ahead 1-year-holding-period excess returns on 2- to 5-year bonds (in rows). Models examined include the macro-yields models using different macroeconomic datasets: the macro-yields, FRED-yields, and LN-yields; their non-Gaussian counterparts: the nonG$_{\text{MY}}$, nonG$_{\text{FRED-Y}}$ and nonG$_{\text{LN-Y}}$; and the predictive regressions on prevailing factors: the CP factor, LN factor, and their combination.

Despite the large dataset appearing capable of improving short-range forecasting, it exhibits significantly inferior performance in 1-year-ahead bond-return

forecasting both in (Table 4.9) and out of sample (Table 4.10). Combining the results with previous analyses, one can recognise the potential of employing different datasets to forecast at different horizons. Rapid advances in ultra-high dimensional copulas afford greater flexibility in modelling large macroeconomic datasets and extracting useful information from them, providing possible directions for future research effort.

Table 4.10: Forecasting excess bond returns using alternative datasets

Out-of-Sample Predictability of Excess Bond Returns (rMSFE)						
Maturity	Macro-yields	nonG_{MY}	FRED-yields	$\text{nonG}_{\text{FRED-Y}}$	LN-yields	$\text{nonG}_{\text{LN-Y}}$
2y	1.37	0.70	1.58	1.08	1.62	1.06
3y	1.30	0.71	1.52	1.07	1.58	1.07
4y	1.27	0.71	1.45	1.04	1.53	1.05
5y	1.28	0.74	1.47	1.06	1.48	1.07

Note. This table reports the out-of-sample rMSFE with respect to the expectations hypothesis for the 1-year-ahead 1-year-holding-period excess returns on 2- to 5-year bonds (in rows).

4.5.2 Economic Value

Inspired by Sarno, Schneider, and Wagner (2016), this section attempts to assess the economic benefits of the improved bond return predictability in the conventional asset allocation framework. The utility-based approach for portfolio optimisation aims to maximise the expected utility $E[u(V_{t+\tau})]$, which takes account of investors' risk preferences as well as the final portfolio value $V_{t+\tau}$ via utility functions u. Among the most-commonly employed ones is the constant coefficient of relative risk aversion (CRRA) utility function, which is widely noted for its scale-invariant feature. Under the CRRA preferences, consider a bond investor allocating between a risk-free bond y_t^τ, whose maturity matches the holding period τ, and a riskier longer-term bond; the expected utility of the end-of-horizon portfolio is

$$E_t[U(V_{t+\tau})] = E_t[\frac{V_{t+\tau}^{1-C}}{1-C}], \tag{4.22}$$

where C denotes the constant relative risk aversion so that $C = -xU''(x)/U'(x)$. In the two-asset allocation problem, let w_t be the initial wealth invested in the risky asset with price B_t and the rest of the amount $1 - w_t$ in the risk-free bond B_t^f; the gross portfolio return over the period is

$$V_{t+\tau}/V_t = 1 + R_{t,t+\tau}^p = 1 + w_t \cdot R_{t,t+\tau}^b + (1 - w_t) \cdot R_{t,t+\tau}^f, \tag{4.23}$$

where $R^b_{t,t+\tau}$ and $R^f_{t,t+\tau}$ refer to the τ-holding period returns of the risky and risk-free bonds, respectively. Substituting logarithmic returns $r^p_{t,t+\tau}$ for arithmetic returns $R^p_{t,t+\tau}$ yields

$$E_t[U(V_{t+\tau})] = \frac{V_t^{1-C}}{1-C} E[(1+R^p_{t,t+\tau})^{1-C}] = \frac{V_t^{1-C}}{1-C} E_t[\exp\{(1-C)r^p_{t,t+\tau}\}]. \quad (4.24)$$

In the continuous case, the dynamics of the portfolio value can be expressed as

$$\frac{dV_t}{V_t} = w_t \cdot \frac{dB_t}{B_t} + (1-w_t) \cdot \frac{dB^f_t}{B^f_t}, \quad (4.25)$$

which is the continuous-time version of Equation (4.23). Given the standard assumption of log-normally distributed returns, applying Ito's lemma and setting $dt = 1$ yields

$$\frac{dV_t}{V_t} = w_t \cdot (r^B_{t,t+1} + \frac{1}{2}\sigma^2_B) + (1-w_t) \cdot r^f_{t,t+1}, \quad (4.26)$$

where $r^B_{t,t+1} = d\log B_t$ and $r^f_{t,t+1} = d\log B^f_t$ correspond to the logarithmic returns and σ^2_B to the conditional variance of risky bonds. Several more stochastic calculations give the log portfolio return

$$r^p_{t,t+1} = d\log V_t = \frac{dV_t}{V_t} - \frac{1}{2}(\frac{dV_t}{V_t})^2 \quad (4.27)$$

$$= r^f_{t,t+1} + w_t(r^B_{t,t+1} - r^f_{t,t+1}) + \frac{1}{2}w_t(1-w_t)\sigma^2_B. \quad (4.28)$$

The closed-form analytical expressions for the expected power utility can be obtained by employing the moment generating function of the log-normal distribution

$$\begin{aligned} E_t[U(V_{t+1})] &= \frac{V_t^{1-C}}{1-C} E_t[\exp\{(1-C)r^p_{t,t+1}\}] \\ &= V_t^{1-C}\{[r^f_{t,t+1} + w_t(E_t[r^B_{t,t+1}] - r^f_{t,t+1})] + \frac{1}{2}w_t(1-Cw_t)\sigma^2_B\}. \end{aligned} \quad (4.29)$$

Taking the first-order conditions provides the optimal weightings

$$w_t = \frac{E_t[r^B_{t,t+1}] - r^f_{t,t+1} + \frac{1}{2}\sigma^2_B}{C\sigma^2_B}, \quad (4.30)$$

where the return predictions come into play and an improved prediction has the potential to facilitate the asset allocation decisions w_t. The relative performance of models is assessed in terms of the manipulation-proof performance measure, as

outlined by Goetzmann, Ingersoll, Spiegel, and Welch (2007)

$$\hat{\Theta}_{i,j} = \frac{1}{(1-\rho)\Delta t} \log \left(\frac{1}{T} \sum_{t=1}^{T} [(1+R_t^i)/(1+R_t^j)]^{1-\rho} \right), \tag{4.31}$$

the amount of which can be seen as the continuously compounded annualised risk-adjusted premium generated by model i over that of the benchmark model j. Compared to conventional measures such as the Sharpe ratio, the manipulation-proof performance measure is more impervious to intentional performance manipulation by fund managers. Further, a risk aversion parameter $\rho = 3$ is chosen following common practice and a maximum leverage exposure of 100% is imposed to inhibit extreme investment decisions.

To be consistent with Section 4.4.2 that examines the predictability of 1-year-holding period returns at 1-year-ahead horizon, this specific analysis evaluates the portfolio performance assuming a 1-year investment horizon, in which case the 1-year bond can be taken as the risk-free bond and the 2-5-year bonds as risky assets. Taking the predictions $r_{t,t+\tau}^{T,\mathrm{M}}$ for bonds of different maturities T=$\{$2y, 3y, 4y, 5y$\}$ from various models M=$\{$OY, MY, nonG$_{\mathrm{OY}}$, nonG$_{\mathrm{MY}}$, CP, LN, LNCP$\}$ as given, the asset allocation is implemented utilising standard tools as described in Equations (4.22–4.30), based on which the economic significance is gauged by the manipulation-proof performance measure. In doing so, the prediction of bond returns and the optimal allocation of assets are treated as relatively separate processes, meaning that the same return-generating process does not necessarily need to be adopted for the return prediction and the portfolio construction.

As indicated by Table 4.11, the enhanced predictability of longer-term bond returns using the non-Gaussian macro-yields model reaps economic gains in an asset allocation context in most cases, which become more pronounced with longer-term bond investing. The non-Gaussian macro-yields model results in positive risk-adjusted premium over the Gaussian scheme and alternative approaches except the LN factor. The relative performance with respect to the expectation hypothesis in terms of statistical predictive accuracy and economic value has received some attention in literature (Campbell and Thompson, 2007; Della Corte et al., 2008; Thornton and Valente, 2012). As suggested by the right panel of Table 4.11, only the non-Gaussian macro-yields model and the LN factor consistently deliver positive premium in excess of the expectations hypothesis across maturities.

Table 4.11: Manipulation-proof performance measure

nonG_{MY}/Model	2y	3y	4y	5y	Model/EH	2y	3y	4y	5y
nonG_{MY}/OY	67.70	121.10	146.18	152.70	OY/EH	-64.32	-110.74	-108.05	-60.67
nonG_{MY}/MY	7.89	22.10	52.44	90.19	MY/EH	-2.25	-6.85	-16.22	-16.27
nonG_{MY}/nonG_{OY}	46.71	67.32	77.92	86.68	nonG_{OY}/EH	-43.71	-57.52	-41.71	1.58
nonG_{MY}/EH	7.23	24.53	59.34	115.31	nonG_{MY}/EH	7.23	24.53	59.34	115.31
nonG_{MY}/CP	68.27	138.79	200.75	249.89	CP/EH	-66.00	-134.83	-178.78	-185.15
nonG_{MY}/LN	-23.86	-43.18	-56.97	-68.92	LN/EH	30.12	63.29	107.85	168.45
nonG_{MY}/LNCP	0.87	9.95	16.62	10.49	LNCP/EH	4.34	5.41	27.13	86.02

Note. This table reports the manipulation-proof performance measures (in basis points) of optimal portfolios that allocate between the 1-year risk-free bond and bonds of different maturities T={2y, 3y, 4y, 5y} (in columns) based on the predictions produced by different models M={OY, MY, nonG_{OY}, nonG_{MY}, CP, LN, LNCP} (in rows). The left panel examines the relative performance of the non-Gaussian macro-yields model with respect to the alternatives while the right panel compares each of the models examined with the expectations hypothesis. A positive value for $\text{Model}_{\text{i}}/\text{Model}_{\text{j}}$ indicates better risk-adjusted performance of Model_{i} relative to Model_{j}.

In the context of Markowitz mean-variance portfolio construction that seeks minimal variance for a target expected rate of return μ^*, the optimal weights are obtained by solving a simple constrained quadratic programming problem

$$\text{minimise} \quad \frac{1}{2} w^T \Sigma w, \tag{4.32}$$

$$\text{subject to} \quad w^T \mu = \mu^*, \mathbb{1}^T w = 1. \tag{4.33}$$

With the inclusion of a risk-fee asset r_f, the weight constraints can be expressed as $w^T \mu + w_f r_f = \mu^*, \mathbb{1}^T w + w_f = 1$, where w and μ refer to the weights and returns for the risky assets. Applying Karush–Kuhn–Tucker conditions gives the optimal weighting for risky assets

$$w = \lambda \Sigma^{-1} (\mu - r_f \mathbb{1}), \tag{4.34}$$

where λ is the Karush–Kuhn–Tucker multiplier concerning the return constraint. Solving the equality conditions gives $\lambda = \frac{(\mu^* - r_f)}{(\mu - r_f \mathbb{1})^T \Sigma^{-1} (\mu - r_f \mathbb{1})}$, yielding the optimal weighting for the risky assets

$$w = \frac{(\mu^* - r_f) \Sigma^{-1} (\mu - r_f \mathbb{1})}{(\mu - r_f \mathbb{1})^T \Sigma^{-1} (\mu - r_f \mathbb{1})}. \tag{4.35}$$

No clear pattern emerges for model comparison in a mean-variance environment, based on which the implications of incorporating macroeconomic variables or employing non-Gaussian approaches cannot be clearly assessed. Although this might simply indicate the departure of economic significance from statistical predictability, the phenomenon is also likely to be due to the inappropriate use of the mean-variance framework, which assumes normal returns or quadratic utility. As

illustrated in Table 4.12, the same predictive values are better exploited by the power utility analysis than the mean-variance approach in terms of enhancing asset allocation decisions, which can be seen from the positive values in 39 out of 40 cases, especially for longer-term bonds. Thus, in this section, the economic value is evaluated primarily in terms of optimal portfolios constructed under CRRA preferences (Table 4.11), suggesting the potential of the non-Gaussian macro-yields model in achieving economic benefits as well as predictive accuracy.

Table 4.12: Power utility vs quadratic utility

Power/Quadratic	2y	3y	4y	5y
OY	13.90	-6.08	21.02	23.51
MY	144.03	251.12	302.37	328.24
nonG_{OY}	48.67	74.39	123.54	161.40
nonG_{MY}	120.83	198.64	289.10	320.47
EH	122.72	209.00	279.25	264.80
CP	28.29	30.22	42.93	28.46
LN	165.64	283.54	375.89	411.13
LNCP	118.89	169.75	229.67	263.14

Note. This table reports the manipulation-proof performance measures (in basis points) of optimal portfolios that allocate between the 1-year risk-free bond and bonds of different maturities T={2y, 3y, 4y, 5y} (in columns) using different utility functions. A positive value for $\text{Model}_{\text{p}}/\text{Model}_{\text{q}}$ indicates better risk-adjusted performance generated by employing power utility function relative to quadratic utility function.

4.6 Conclusion

This research explores the informational content of macroeconomic variables by means of copulas in modelling and forecasting the yield curve and excess bond returns. Collectively, these results underscore the crucial importance of incorporating macroeconomic variables and accounting for non-Gaussianity inherent in yield distributions. This research lays the groundwork for deploying copula techniques in the context of macro-finance modelling. Advances over the past decade in high-dimensional copula models and applications have resulted in a wider range of copula tools being available, which accommodate more general and realistic structures. Considerably more work will be necessary in order to pursue more appropriate copula models and marginal density estimators, as well as bandwidth selection procedures that achieve superior performance-complexity trade-offs.

4.7 Appendices

4.A Estimation of the Inversion Copula

Besides the greater flexibility discussed above, the copula method has the added advantage of faster and simpler estimation. In the traditional maximum-likelihood estimation, the log-likelihood function based on n observations $Y = (y_1^t, \cdots, y_n^t)_{t=1}^T$ can be expressed in terms of copula and marginal densities

$$\hat{\theta} = \arg\max_{\theta}\{l(\theta) = \sum_{t=1}^{T} \ln c(F_1(y_1^t), \ldots, F_n(y_n^t)) + \sum_{t=1}^{T}\sum_{i=1}^{n} \ln f_i(y_i^t)\} \tag{4.36}$$

and is maximised over the whole parameter space Θ. Simultaneously estimating the parameters of the margins and copula involves significant computational challenges, particularly for high dimensional cases. Exploiting the copula-marginal separation, the inference function for margins (IFM) method (Joe, 1997) implements a two-stage procedure, performing calibrations of the marginal parameters $\theta_1, \cdots, \theta_d$ and the copula parameters θ_C sequentially

$$\hat{\theta} = \arg\max_{\theta}\{l(\theta) = \sum_{t=1}^{T} \ln c(F_1(y_1^t; \theta_1), \ldots, F_n(y_n^t, \theta_n); \theta_C) + \sum_{t=1}^{T}\sum_{i=1}^{n} \ln f_i(y_i^t; \theta_i)\}. \tag{4.37}$$

The marginal parameters $\theta_1, \cdots, \theta_d$ are first obtained by employing the maximum likelihood estimation $\hat{\theta}_i = \arg\max_{\theta_i} \{\sum_{t=1}^{T} \ln f_i(y_i^t; \theta_i)\}$, based on which the copula parameters θ_C are estimated to yield the whole IFM estimator $\hat{\theta} = (\hat{\theta_1}, \cdots, \hat{\theta_n}, \hat{\theta_C})$

$$\hat{\theta_C} = \arg\max_{\theta_C}\{\sum_{t=1}^{T} \ln c(F_1(y_1^t; \hat{\theta_1}), \ldots, F_n(y_n^t, \hat{\theta_n}); \theta_C)\}. \tag{4.38}$$

In a two-stage procedure similar to the inference function for margins, the canonical maximum likelihood method first transforms the original observations $Y = (y_1^t, \cdots, y_n^t)_{t=1}^T$ into uniformly distributed pseudo-observations $\hat{U} = (\hat{u}_1^t, \cdots, \hat{u}_n^t)_{t=1}^T = (\hat{F}_1(y_1^t), \cdots, \hat{F}_n(y_n^t))_{t=1}^T$, which are then fed into the optimiser to maximise the copula log-likelihood

$$\hat{\theta_C} = \arg\max_{\theta_C}\{\sum_{t=1}^{T} \ln c(\hat{u}_1^t, \ldots, \hat{u}_n^t; \theta_C)\}. \tag{4.39}$$

4.B Marginal Densities

Compelling motivations for scrutinising the normality assumption of Gaussian macro-yields models also arise from the severe performance deterioration when implemented in the 1970:M1–2016:M12 sample that extends to the post-global financial crisis period. The decline in performance is largely attributable to the yields' standing at incredibly low levels since the crisis, which establishes positive asymmetry and a density spike near zero. In sharp contrast, macroeconomic variables exhibit significant fluctuations, as shown in the growth rate of the consumer price index, a widely-used measure of inflation, and that of the industrial production index, a representative indicator of real economic activity (Figure 4.7). This discrepancy poses significant challenges for the linear-Gaussian structure that used to yield satisfactory results before the crisis, and thus leads to inferior performance over the extended period.

It should be noted that the non-Gaussian features are not exclusive to the yield distributions, and have appeared in various macroeconomic series. Accounting more fully for the non-Gaussian nature of macroeconomic variables as well as yields maintains a competitive advantage over the Gaussian setting, although in general the improvement is not as remarkable as that generated by the current specification that focuses on non-Gaussianity in yields. Properties provided by applying non-Gaussianity only to yield variables may be preferable in such a case where yield forecasting is of ultimate interest. This chapter does not intend to optimise this choice, owing to its satisfactory performance in practice, but simultaneously acknowledges the possibility that the current representation is not guaranteed to be desirable for other purposes, for example in forecasting macroeconomic variables.

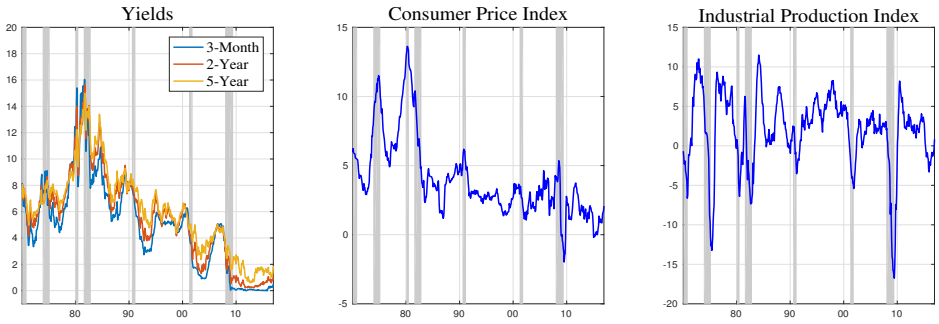

Figure 4.7: Time series of yields and macroeconomic indicators. This figure displays time series of yields (3-month, 2-year, 5-year) and two representative macroeconomic indicators, the growth rates of consumer price index and industrial production index, over the timespan 1970:M1–2016:M12. The evolution of business cycles is indicated by the grey-shaded areas that correspond to economic contractions identified by NBER.

4.C Normality Test

Table 4.13: Normality test for yields

Period/Maturity	3m	1y	2y	3y	4y	5y
1970:M1–1987:M6	0.93***	0.94***	0.93***	0.92***	0.92***	0.92***
1970:M1–2008:M12	0.95***	0.96***	0.97***	0.97***	0.97***	0.96***
1970:M1–2016:M12	0.95***	0.96***	0.96***	0.97***	0.97***	0.97***

Note. This table reports the p-values of the Shapiro-Wilk (Shapiro and Wilk, 1965) or Shapiro-Francia (Shapiro and Francia, 1972) tests for normality for each maturity (in columns) over expanding time periods: 1970:M1–1987:M6, 1970:M1–2008:M12, and 1970:M1–2016:M12 (in rows). Rejection of the null hypothesis at significance level 1%, implying departure from normality, is marked with ***.

4.D Test for Superior Predictive Accuracy

The test for predictive accuracy is implemented using the bootstrap data snooper in Kevin Sheppard's MFE toolbox that carries out the reality check of White (2000) and the superior predictive accuracy test of Hansen (2005). Rejection of the null hypothesis indicates that the examined model generates significantly smaller loss than the benchmark, which in this specific context means that the outperformance of the non-Gaussian scheme relative to its Gaussian counterpart or the no-predictability benchmark is statistically significant. Such testing techniques are developed with the purpose of mitigating the data-snooping biases in statistical inference and model selection. When comparing multiple models or strategies M with the benchmark, the reality check constructs the following adjusted statistics

$$\bar{V} = \max_{i=1,\dots,M} \{\sqrt{N}\bar{f}_i\}, \tag{4.40}$$

$$\bar{V}_B^* = \max_{i=1,\dots,M} \{\sqrt{N}(\bar{f}_{i,j}^* - \bar{f}_i)\}, j = 1, \dots, B, \tag{4.41}$$

where $\bar{f}_{i,j}^*$ is the average of j^{th} bootstrap resampling of i^{th} model with replacement, i.e. $\bar{f}_{i,j}^* = n^{-1}\sum_k f_{i,j,k}$. Comparing $\bar{V}$ and the distribution of $\bar{V}_B^*$ yields the reality check bias-adjusted p-values.

This chapter performs pairwise comparisons among models, in which case only one model $M = 1$ is examined against its counterpart (as a benchmark) at each time. The superior predictive ability test proposed by Hansen (2005) boosts the power of the test by accounting for the variance estimates of the residuals. It is noteworthy that due to the use of bootstrapping techniques, the test does not necessarily yield the same estimates of statistics at each implementation, and thus leads to slightly different significance levels but qualitatively similar conclusions.

For more details on the predictive accuracy test, please refer to White (2000) and Hansen (2005).

4.E Smoothed Squared Errors over Time

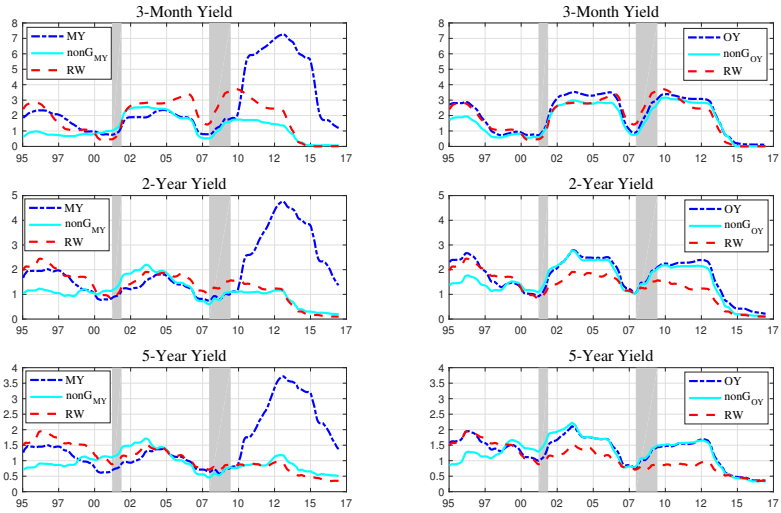

Figure 4.8: Forecasting the yield curve over time (smoothed squared errors). This figure displays the smoothed squared errors of 1-year-ahead forecasts for 3m-(upper), 2y-(middle), and 5y-yields (lower) over a 5-year rolling window, with the end dates of rolling windows reported along the x-axis. The dash-dotted lines represent the original Gaussian models (MY, OY) while the continuous lines plot their non-Gaussian counterparts (nonG_{MY}, nonG_{OY}), with the dashed lines reporting the random walk benchmark (RW).

Following the idea of the fluctuation test proposed by Giacomini and Rossi (2010), which examines the local relative performance and its time-varying patterns, Figure 4.10 reports the out-of-sample Diebold–Mariano statistics of 1-year-ahead (12-step-ahead) forecasts for 1-year-holding-period excess returns on 5-year yield, computed over a 5-year rolling window to be in line with Giacomini and Rossi (2010) and Coroneo, Giannone, and Modugno (2016). Positive values refer to lower MSFEs of the non-Gaussian macro-yields model, and thus indicate outperformance with respect to the Gaussian macro-yields model (upper panel) and the expectations hypothesis (lower panel), respectively[3]. The local relative performance in Figure 4.10 and the long-term average performance in Table 4.5 together indicate that the

[3]The critical value for the fluctuation test using a fixed or rolling estimation scheme is reported by Giacomini and Rossi (2010). For a recursive scheme, West (1996) examines the case where the uncertainty of parameter estimation is also taken account of. Though it could be obtained in principal, the critical value in the case of recursive scheme is not readily available as those for the fixed or rolling scheme and is beyond the scope of this analysis. Smoothed squared errors using rolling windows as reported by Coroneo, Giannone, and Modugno (2016) are displayed in Appendix 4.E for a direct comparison with their results.

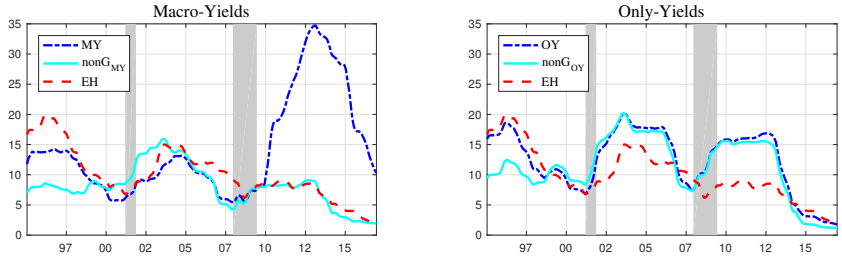

Figure 4.9: Forecasting excess bond returns over time (smoothed squared errors). This figure displays the smoothed squared errors of 1-year-ahead forecasts for 1-year-holding-period excess returns over a 5-year rolling window, with the end dates of rolling windows reported along the x-axis. The dash-dotted lines represent the original Gaussian models (MY, OY) while the continuous lines plot their non-Gaussian counterparts (nonG_{MY}, nonG_{OY}), with the dashed lines reporting the expectations hypothesis benchmark (EH). The evolution of business cycles is indicated by the grey-shaded areas that correspond to economic contractions identified by NBER.

improved forecast accuracy afforded by the proposed copula is neither specific to a particular time-period nor subject to a one-time reversal. In particular, the non-Gaussian macro-yields model produces plausible results in recent periods, which aligns with its average performance over the whole sample period.

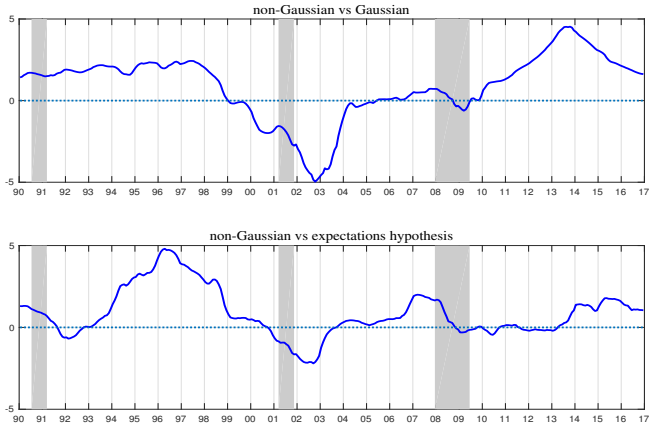

Figure 4.10: Local relative forecasting performance. This figure displays the Diebold–Mariano test statistics of 1-year-ahead forecasts for 1-year-holding-period excess returns on 5-year bonds over a 5-year rolling window, with the end dates of rolling windows reported along the x-axis. The upper panel examines the loss-differential of the non-Gaussian macro-yields model with respect to the Gaussian macro-yields model, while the lower panel examines that with respect to the expectations hypothesis.

4.F Rolling Window Analysis

Table 4.14: Forecasting the yield curve using rolling windows

	Expanding-Window						Rolling-Window					
	3m	1y	2y	3y	4y	5y	3m	1y	2y	3y	4y	5y
h=1	1.026	0.967	1.078	1.043	1.046	1.020	0.959	1.073	1.081	1.108	1.085	1.053
h=3	0.862	0.865	1.006	0.985	1.020	1.012	0.726	0.873	0.961	0.951	0.987	1.001
h=6	0.731	0.804	0.925	0.934	0.975	0.977	0.703	0.869	0.978	0.981	1.014	1.029
h=12	0.588	0.673	0.762	0.800	0.838	0.857	0.771	0.923	1.001	1.039	1.067	1.088
h=24	0.640	0.743	0.835	0.898	0.954	1.027	1.057	1.194	1.330	1.419	1.485	1.565

Note. This table reports the rMSFE of the non-Gaussian macro-yields model with respect to the random walk for 3m-, 1y-, 2y-, 3y-, 4y-, and 5y-yields (in columns) over horizons of h=1, 3, 6, 12, and 24 months (in rows) using a recursive scheme (left) and using a rolling scheme (right). Forecasts are constructed using observations from 1970:M1 and are evaluated over 1990:M1–2016:M12 as the forecast origin evolves.

Table 4.15: Forecasting excess bond returns using rolling windows

	Expanding-Window		Rolling-Window	
	non-Gaussian MY	non-Gaussian OY	non-Gaussian MY	non-Gaussian OY
2y	0.70	1.08	0.95	1.15
3y	0.71	1.05	0.94	1.11
4y	0.71	1.01	0.92	1.06
5y	0.74	1.02	0.94	1.05

Note. This table reports the rMSFE of the non-Gaussian macro-yields model with respect to the expectations hypothesis for excess bond returns using a recursive scheme (left) and using a rolling scheme (right). Forecasts are constructed using observations from 1970:M1 and are evaluated over 1990:M1–2016:M12 as the forecast origin evolves.

4.G Results Using a Single Fixed Bandwidth

This subsection reports the results of the previous analyses using a single fixed bandwidth across maturities and over time. A series of optimal bandwidths is first generated utilising a recursive scheme. Averaging the optimal estimates over time and across maturities generates a kernel bandwidth of 0.1503 using the golden-section search and 0.1556 when a selection set of evenly-spaced points in the interval [0.05, 0.8] with increments of 0.05 is pre-specified. Accordingly, this section re-conducts the examination of predictability by employing a single fixed bandwidth of 0.15. A comparison with Table 4.2 and Table 4.5 detects only small differences. While the fixed bandwidth slightly outperforms in terms of yield forecasting, the time-varying bandwidth displays greater in-sample and out-of-sample bond return predictability. In either case, the conclusions regarding the benefits of adopting the non-Gaussian scheme and incorporating macroeconomic variables remain unchanged.

Table 4.16: Forecasting the yield curve with a single fixed bandwidth

	non-Gaussian OY						non-Gaussian MY					
	3m	1y	2y	3y	4y	5y	3m	1y	2y	3y	4y	5y
h=1	0.822	0.971	1.053	1.022	1.018	0.996	1.026	0.977	1.075	1.037	1.048	1.022
h=3	0.773	0.980	1.070	1.041	1.033	1.029	0.869	0.873	0.996	0.972	1.003	0.999
h=6	0.815	1.032	1.104	1.084	1.074	1.077	0.742	0.818	0.920	0.918	0.951	0.956
h=12	0.884	1.053	1.132	1.153	1.158	1.167	0.623	0.714	0.787	0.809	0.835	0.848
h=24	0.964	1.111	1.253	1.363	1.453	1.543	0.681	0.778	0.860	0.926	0.973	1.035

Note. This table reports the rMSFE of the non-Gaussian only-yields model (left) and non-Gaussian macro-yields model (right) with respect to the random walk, with a single fixed bandwidth of 0.15 for the kernel density estimator.

Table 4.17: Forecasting excess bond returns with a single fixed bandwidth

In-Sample R^2							
Maturity	OY	non-Gaussian OY	MY	non-Gaussian MY	CP	LN	LN+CP
2y	12.35%	15.24%	40.36%	40.84%	15.58%	22.26%	29.62%
3y	13.21%	17.23%	38.35%	37.93%	18.09%	24.78%	33.52%
4y	14.70%	19.75%	34.66%	35.13%	21.05%	25.22%	36.07%
5y	15.89%	22.51%	31.37%	34.74%	18.69%	25.10%	34.22%
Out-of-Sample rMSFE							
Maturity	OY	non-Gaussian OY	MY	non-Gaussian MY	CP	LN	LN+CP
2y	1.23	1.09	1.37	0.74	1.32	1.22	0.96
3y	1.23	1.06	1.30	0.74	1.40	1.07	0.90
4y	1.18	1.02	1.27	0.72	1.41	0.97	0.84
5y	1.13	1.02	1.28	0.74	1.40	0.90	0.82

Note. This table reports the in-sample R^2 and out-of-sample rMSFE with respect to the expectations hypothesis for excess bond returns, using a single fixed bandwidth of 0.15.

4.H Different Choices of Bandwidth

The copula framework is readily extensible to explore various sets of variables and marginal density estimators, boosting its potential in terms of enhancing forecasting for different maturities and forecast horizons. As indicated previously, the selection of the smoothing parameter in kernel estimation, that is the bandwidth, is crucial to characterising the marginal distributions accurately, and thus has a substantial influence on the resulting performance. Intending to promote out-of-sample forecasting rather than in-sample fitting, this section attempts to compare and select the bandwidth in terms of predictability. However, it is beyond the scope of this chapter to establish rigorous criteria and procedures for general selection problems.

One might speculate on the possibility of developing a selection procedure that enables the bandwidth to vary for each forecast origin and to vary by forecast horizons. However, the primary concern that arises is how to develop selection cri-

teria in terms of out-of-sample performance. To deal with this matter, at each point in time, the information set up to time T_e is further divided into a sub-estimation-period $[t_0, t_e]$, over which the models are fitted using different bandwidths, and a sub-evaluation-period $[t_e, T_e - \tau_h]$, over which the prediction performance of each bandwidth is assessed. Using the optimal bandwidth selected for horizon h, models are re-estimated based on the information set $[t_0, T_e]$, with which the forecasts are then generated.

To ensure a certain number of observations for estimation, 1985:M1 is employed as the partition of the sub-estimation and sub-evaluation periods. Taking T_e = 1990M1 as an illustration, yield forecasts are generated using data from t_0 = 1970:M1–1985:M1 and evaluated over the subsample [1985:M1, 1989:M12] for 1-month forecasts, and [1985:M1, 1989:M9] for 3-month forecasts, ..., to [1985:M1, 1988:M1] for 24-month forecasts. The procedure is repeated for each of the evenly spaced points in the interval [0.05, 0.8] with increments of 0.05. The bandwidth that delivers the best h-step forecasts for most maturities is selected for the 1990:M1 evaluation point. To avoid introducing excessive variation, the forecast horizons are divided into three groups, short-term (h=1), medium-term (h=3, 6, 12), and long-term (h=24), while each group shares a common bandwidth that yields superior average performance.

Table 4.18 reports the out-of-sample rMSFEs for yields utilising bandwidths that vary across forecast origins and vary by forecast horizons. The non-Gaussian macro-yields model maintains its competitive advantage compared to the vanilla macro-yields model, with slight deterioration in the medium-term forecasts and significant improvements in the long-term forecasts. The shortest-range forecasting, however, remains a challenge, emphasising the importance of exploring other information resources. The results are more stable and consistent in the case of bond return forecasting, arriving at similar conclusions.

Table 4.18: Varying bandwidths for different horizons

	3m	1y	2y	3y	4y	5y
h=1	1.078	1.036	1.117	1.070	1.074	1.040
h=3	0.885	0.892	1.002	0.975	0.996	0.987
h=6	0.806	0.894	0.976	0.960	0.969	0.954
h=12	0.669	0.757	0.807	0.811	0.820	0.813
h=24	0.567	0.622	0.658	0.706	0.746	0.794

Note. This table reports the rMSFE for yields of the non-Gaussian macro-yields model using a varying bandwidth across horizons, which is selected according to the out-of-sample predictive performance and is updated recursively over time.

Chapter 5

Conclusions and Future Work

5.1 Conclusions

This thesis has made some progress in terms of enhancing the understanding of the underlying determinants of stock and bond returns, with particular focus on their predictive content.

Chapter 2 enhances the cross-sectional stock return predictability by exploiting the informational content of variance swaps. Compared to parametric stochastic models, the non-parametric principal component approach adopted in this chapter enables extraction of information not only on the underlying states of the asset price process (e.g. the instantaneous variance, the long-run variance, and the jump) but also on the associated risk premia which reflect market perceptions. The incorporation of variance swap components results in improved forecasts by contributing further insights about the overall market conditions beyond those provided by the benchmark factors.

Chapter 3 reveals substantial deteriorations in the predictive power of macroeconomic variables and instabilities in the macro-yields dependency after the global financial crisis in the context of the linear-Gaussian dynamic factor representation. This chapter suggests incapability of the linear-Gaussian factor framework to capture the post-crisis evolutionary patterns of yields and to characterise the post-crisis macro-finance interactions.

Chapter 4 develops an effective modelling strategy that deals considerably better than the linear-Gaussian DFM with the post-2008 sample. This chapter proposes a non-Gaussian macro-yields model to enable more complex characterisations of yield distributions and macro-yields interdependence, particularly to account for the asymmetry in yield distributions. This non-Gaussian model achieves markedly superior performance compared to the Gaussian macro-yields model and the no-

predictability benchmarks in terms of both in-sample and out-of-sample forecasting for the yield curve and excess bond returns. Furthermore, this copula-based approach affords a technically convenient and extensible means of accommodating high-dimensional macroeconomic datasets and the growing complexity of post-crisis yield movements, which facilitates further investigation into their practical implications for government bond modelling and forecasting.

Collectively, while the macroeconomic panel does convey additional insights into the future development of yields, the manner in which macroeconomic information is exploited is pivotal in introducing macroeconomic perspectives into yield modelling and forecasting. Chapter 4 provides a useful avenue for exploring the predictive value of high-dimensional macroeconomic datasets in a low interest rate environment.

5.2 Future Work

Once deemed unthinkable, the negative-yielding bond is gradually becoming the norm, predominantly in the Eurozone and Japan, owing to policymakers' aggressive adoption of conventional and unconventional monetary policies in order to nurture economic growth in the aftermath of the global financial crisis. The market has recently witnessed Germany's 10-year Bund yield turning negative, the longest-dated Treasury yield plunging to record lows, and negative-yielding bonds constituting nearly one-third of the global sovereign debt market, along with exacerbated concerns regarding the health of the global economy. The negative-yielding bond still has its place in the portfolio of loss-averse investors who hold pessimistic beliefs and view government bonds as relatively safe-haven assets, as well as those who take account of the possibility of currency appreciation or deflation and gauge profit in real rather than nominal terms. This also holds true for institutional investors who purchase bonds for specific purposes, such as for collateral use and asset-liability management.

Chapter 3 has revealed troubling issues with the conventional macro-finance DFM in handling the persistently low interest rates. Unprecedented monetary stimulus by means of aggressive rate cuts and quantitative easing, along with slower productivity growth and changing demographics that reshape saving behaviour are causing economic pain throughout the world, which is particularly acute in the Eurozone and Japan. This leaves abundant room for future investigations into the post-crisis performance of macro-finance term structure models in these sovereign bond markets and into the stability of the model-implied macro-financial linkages.

Chapter 4 explores the potential of the copula-based non-Gaussian macro-

finance model in delivering robust performance before and after the crisis. Considering the rising life expectancy, heightened geopolitical risks and lacklustre economic growth, amongst other influencing factors, the possibility cannot be safely ruled out that U.S. Treasury yields might remain stuck in low gear if the Federal Reserve initiates a new round of rate-cutting and quantitative easing. Given the current global landscape being bombarded with such uncertainty, the approach proposed here facilitates greater flexibility by allowing the yields to either tip into negative territory or to recover to pre-crisis levels. Moreover, the approach has demonstrated promising performance regarding U.S. Treasury securities by accounting for post-crisis features that cannot be adequately addressed by the traditional Gaussian representation. The underlying mechanism generally applies to other sovereign bond markets, and is likely to be even more appealing for those already suffering from negative rates, such as Switzerland, Germany, the Netherlands, Japan, and Denmark.

As discussed in Chapter 2, the multi-period asset pricing models take into account consumption and investment in future periods when determining the current cross-sectional risk-return trade-off. Investors require less compensation for assets that hedge against future deterioration of investment opportunities, while expecting premiums to be higher for those that co-vary with solid economic growth and declining marginal utility. The literature attempts to decipher clues regarding state variables that describe the development of the investment opportunity set from widely-recognised time-series predictors of aggregate stock returns and economic activity, such as the variance and jump-related components presented in Chapter 2 and macro-financial factors as examined by other studies (Brennan et al., 2004; Cochrane and Piazzesi, 2005; Petkova, 2006; Lettau and Wachter, 2011; Koijen et al., 2017). A growing body of research is seeking to simultaneously model bonds and stocks, at the aggregate level or in a cross-sectional context (Boons, 2016; Koijen et al., 2017). In future investigations, it would be interesting to examine the information content of the yield and macro factors discussed in Chapter 3 in the context of forecasting the cross-sectional stock returns and advancing the unified equity-bond modelling framework, as well as exploring the potential of the copula-based approach considered in Chapter 4 for tackling the asymmetry and heavy-tailedness in the return distributions of other asset classes.

Bibliography

Kjersti Aas, Claudia Czado, Arnoldo Frigessi, and Henrik Bakken. Pair-copula constructions of multiple dependence. *Insurance: Mathematics and Economics*, 44(2):182–198, 2009.

Knut Are Aastveit, Andrea Carriero, Todd E Clark, and Massimiliano Marcellino. Have standard VARs remained stable since the crisis? *Journal of Applied Econometrics*, 32(5):931–951, 2017.

Abhay Abhyankar and Angelica Gonzalez. News and the cross-section of expected corporate bond returns. *Journal of Banking & Finance*, 33(6):996–1004, 2009.

Tobias Adrian and Joshua Rosenberg. Stock returns and volatility: Pricing the short-run and long-run components of market risk. *The Journal of Finance*, 63 (6):2997–3030, 2008.

Yacine Ait-Sahalia, Mustafa Karaman, and Loriano Mancini. The term structure of variance swaps and risk premia. *Swiss Finance Institute Research Paper*, (18-37), 2018.

Carlos Almeida, Claudia Czado, and Hans Manner. Modeling high-dimensional time-varying dependence using dynamic D-vine models. *Applied Stochastic Models in Business and Industry*, 32(5):621–638, 2016.

Carlo Altavilla, Raffaella Giacomini, and Giuseppe Ragusa. Anchoring the yield curve using survey expectations. *Journal of Applied Econometrics*, 32(6):1055–1068, 2017.

Dante Amengual. The term structure of variance risk premia. *Princeton University*, 2009.

Byeong-Je An, Andrew Ang, Turan G Bali, and Nusret Cakici. The joint cross section of stocks and options. *The Journal of Finance*, 69(5):2279–2337, 2014.

Leif Andersen and Jakob Sidenius. Extensions to the gaussian copula: Random recovery and random factor loadings. *Journal of Credit Risk Volume*, 1(1):05, 2004.

Andrew Ang and Monika Piazzesi. A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables. *Journal of Monetary Economics*, 50(4):745–787, 2003.

Andrew Ang, Robert J Hodrick, Yuhang Xing, and Xiaoyan Zhang. The cross-section of volatility and expected returns. *The Journal of Finance*, 61(1):259–299, 2006.

Andrew Ang, Sen Dong, and Monika Piazzesi. No-arbitrage Taylor rules. Technical report, National Bureau of Economic Research, 2007.

Andrew Ang, Jean Boivin, Sen Dong, and Rudy Loo-Kung. Monetary policy shifts and the term structure. *The Review of Economic Studies*, 78(2):429–457, 2011.

Craig F Ansley and Robert Kohn. A geometrical derivation of the fixed interval smoothing algorithm. *Biometrika*, 69(2):486–487, 1982.

Jushan Bai and Serena Ng. Determining the number of factors in approximate factor models. *Econometrica*, 70(1):191–221, 2002.

Gurdip Bakshi and Nikunj Kapadia. Delta-hedged gains and the negative market volatility risk premium. *The Review of Financial Studies*, 16(2):527–566, 2003.

Gurdip Bakshi and Dilip Madan. Spanning and derivative-security valuation. *Journal of Financial Economics*, 55(2):205–238, 2000.

Gurdip Bakshi, Nikunj Kapadia, and Dilip Madan. Stock return characteristics, skew laws, and the differential pricing of individual equity options. *The Review of Financial Studies*, 16(1):101–143, 2003.

Turan G Bali and Scott Murray. Does risk-neutral skewness predict the cross-section of equity option portfolio returns? *Journal of Financial and Quantitative Analysis*, 48(4):1145–1171, 2013.

Turan G Bali, Robert F Engle, and Scott Murray. *Empirical asset pricing: The cross section of stock returns*. John Wiley & Sons, 2016.

Marta Bańbura, Domenico Giannone, and Michele Lenza. Conditional forecasts and scenario analysis with vector autoregressions for large cross-sections. *International Journal of forecasting*, 31(3):739–756, 2015.

Michael D Bauer and Glenn D Rudebusch. Monetary policy expectations at the zero lower bound. *Journal of Money, Credit and Banking*, 48(7):1439–1465, 2016a.

Michael D Bauer and Glenn D Rudebusch. Resolving the spanning puzzle in macro-finance term structure models. *Review of Finance*, 21(2):511–553, 2016b.

Michael D Bauer, Glenn D Rudebusch, and Jing Cynthia Wu. Term premia and inflation uncertainty: Empirical evidence from an international panel dataset: Comment. *American Economic Review*, 104(1):323–37, 2014.

Richard A Berk. *Statistical Learning from a Regression Perspective*. Springer, 2016.

Ben S Bernanke and Jean Boivin. Monetary policy in a data-rich environment. *Journal of Monetary Economics*, 50(3):525–546, 2003.

Ben S Bernanke, Jean Boivin, and Piotr Eliasz. Measuring the effects of monetary policy: A factor-augmented vector autoregressive (FAVAR) approach. *The Quarterly Journal of Economics*, 120(1):387–422, 2005.

Francesco Bianchi, Haroon Mumtaz, and Paolo Surico. The great moderation of the term structure of UK interest rates. *Journal of Monetary Economics*, 56(6): 856–871, 2009.

Ruslan Bikbov and Mikhail Chernov. No-arbitrage macroeconomic determinants of the yield curve. *Journal of Econometrics*, 159(1):166–182, 2010.

Bahar Biller. Copula-based multivariate input models for stochastic simulation. *Operations research*, 57(4):878–892, 2009.

Bahar Biller and Barry L Nelson. Modeling and generating multivariate time-series input processes using a vector autoregressive technique. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 13(3):211–237, 2003.

Tim Bollerslev, George Tauchen, and Hao Zhou. Expected stock returns and variance risk premia. *The Review of Financial Studies*, 22(11):4463–4492, 2009.

Tim Bollerslev, Michael Gibson, and Hao Zhou. Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities. *Journal of Econometrics*, 160(1):235–245, 2011.

Tim Bollerslev, James Marrone, Lai Xu, and Hao Zhou. Stock return predictability and variance risk premia: Statistical inference and international evidence. *Journal of Financial and Quantitative Analysis*, 49(3):633–661, 2014.

Martijn Boons. State variables, macroeconomic activity, and the cross section of individual stocks. *Journal of Financial Economics*, 119(3):489–511, 2016.

Michael J Brennan, Ashley W Wang, and Yihong Xia. Estimation and test of a simple model of intertemporal capital asset pricing. *The Journal of Finance*, 59 (4):1743–1776, 2004.

Mark Broadie, Mikhail Chernov, and Michael Johannes. Model specification and risk premia: Evidence from futures options. *The Journal of Finance*, 62(3):1453–1490, 2007.

Andrea Buraschi, Fabio Trojani, and Andrea Vedolin. Economic uncertainty, disagreement, and credit markets. *Management Science*, 60(5):1281–1296, 2013.

Gonzalo Camba-Mendez. Conditional forecasts on SVAR models using the Kalman filter. *Economics Letters*, 115(3):376–378, 2012.

John Y Campbell and Samuel B Thompson. Predicting excess stock returns out of sample: Can anything beat the historical average? *The Review of Financial Studies*, 21(4):1509–1531, 2007.

Mark M Carhart. On persistence in mutual fund performance. *The Journal of Finance*, 52(1):57–82, 1997.

Peter Carr and Liuren Wu. Leverage effect, volatility feedback, and self-exciting market disruptions. *Journal of Financial and Quantitative Analysis*, 52(5):2119–2156, 2017.

Alberto Caruso and Laura Coroneo. Predicting interest rates in real-time. 2019.

Fousseni Chabi-Yo. Pricing kernels with stochastic skewness and volatility risk. *Management Science*, 58(3):624–640, 2012.

Bo Young Chang, Peter Christoffersen, and Kris Jacobs. Market skewness risk and the cross section of stock returns. *Journal of Financial Economics*, 107(1):46–68, 2013.

Xu Cheng, Zhipeng Liao, and Frank Schorfheide. Shrinkage estimation of high-dimensional factor models with structural instabilities. *The Review of Economic Studies*, 83(4):1511–1543, 2016.

Mikhail Chernov and Philippe Mueller. The term structure of inflation expectations. *Journal of Financial Economics*, 106(2):367–394, 2012.

Jens HE Christensen and Glenn D Rudebusch. Estimating shadow-rate term structure models with near-zero yields. *Journal of Financial Econometrics*, 13(2): 226–259, 2014.

Jens HE Christensen and Glenn D Rudebusch. Modeling yields at the zero lower bound: Are shadow rates the solution? In *Dynamic Factor Models*, pages 75–125. Emerald Group Publishing Limited, 2016.

Jens HE Christensen, Francis X Diebold, and Glenn D Rudebusch. An arbitrage-free generalized Nelson–Siegel term structure model. *The Econometrics Journal*, 12(3):C33–C64, 2009.

Jens HE Christensen, Jose A Lopez, and Glenn D Rudebusch. Inflation expectations and risk premiums in an arbitrage-free model of nominal and real bond yields. *Journal of Money, Credit and Banking*, 42:143–178, 2010.

Jens HE Christensen, Francis X Diebold, and Glenn D Rudebusch. The affine arbitrage-free class of Nelson–Siegel term structure models. *Journal of Econometrics*, 164(1):4–20, 2011.

Peter Christoffersen, Vihang Errunza, Kris Jacobs, and Hugues Langlois. Is the potential for international diversification disappearing? a dynamic copula approach. *The Review of Financial Studies*, 25(12):3711–3751, 2012.

Peter Christoffersen, Kris Jacobs, Xisong Jin, and Hugues Langlois. Dynamic dependence and diversification in corporate credit. *Review of Finance*, 22(2):521–560, 2017.

Anna Cieslak and Pavol Povala. Expected returns in Treasury bonds. *The Review of Financial Studies*, 28(10):2859–2901, 2015.

Todd E Clark and Michael W McCracken. Tests of predictive ability for vector autoregressions used for conditional forecasting. *Journal of Applied Econometrics*, 32(3):533–553, 2017.

John H Cochrane and Monika Piazzesi. Bond risk premia. *American Economic Review*, 95(1):138–160, 2005.

Jennifer Conrad, Robert F Dittmar, and Eric Ghysels. Ex ante skewness and expected stock returns. *The Journal of Finance*, 68(1):85–124, 2013.

Ilan Cooper and Richard Priestley. Time-varying risk premiums and the output gap. *The Review of Financial Studies*, 22(7):2801–2833, 2008.

Laura Coroneo, Ken Nyholm, and Rositsa Vidova-Koleva. How arbitrage-free is the Nelson–Siegel model? *Journal of Empirical Finance*, 18(3):393–407, 2011.

Laura Coroneo, Domenico Giannone, and Michele Modugno. Unspanned macroeconomic factors in the yield curve. *Journal of Business & Economic Statistics*, 34 (3):472–485, 2016.

Drew D Creal and Ruey S Tsay. High dimensional dynamic stochastic copula models. *Journal of Econometrics*, 189(2):335–345, 2015.

Martijn Cremers and David Weinbaum. Deviations from put-call parity and stock return predictability. *Journal of Financial and Quantitative Analysis*, 45(2):335–367, 2010.

Martijn Cremers, Michael Halling, and David Weinbaum. Aggregate jump and volatility risk in the cross-section of stock returns. *The Journal of Finance*, 70 (2):577–614, 2015.

Dean Croushore and Tom Stark. A real-time data set for macroeconomists. *Journal of Econometrics*, 105(1):111–130, 2001.

Michiel De Pooter, Francesco Ravazzolo, and Dick JC Van Dijk. Term structure forecasting using macro factors and forecast combination. *FRB International Finance Discussion Paper*, (993), 2010.

Davide Debortoli, Jordi Galí, and Luca Gambetti. On the empirical (ir) relevance of the zero lower bound constraint. Technical report, National Bureau of Economic Research, 2019.

Pasquale Della Corte, Lucio Sarno, and Daniel L Thornton. The expectation hypothesis of the term structure of very short-term rates: Statistical tests and economic value. *Journal of Financial Economics*, 89(1):158–174, 2008.

Francis X Diebold and Canlin Li. Forecasting the term structure of government bond yields. *Journal of Econometrics*, 130(2):337–364, 2006.

Francis X Diebold and Robert S Mariano. Comparing predictive accuracy. *Journal of Business & economic statistics*, 20(1):134–144, 2002.

Francis X Diebold, Glenn D Rudebusch, and S Boragan Aruoba. The macroeconomy and the yield curve: A dynamic latent factor approach. *Journal of Econometrics*, 131(1-2):309–338, 2006.

George Dotsis, Dimitris Psychoyios, and George Skiadopoulos. An empirical comparison of continuous-time models of implied volatility indices. *Journal of Banking & Finance*, 31(12):3584–3603, 2007.

Itamar Drechsler and Amir Yaron. What's vol got to do with it. *The Review of Financial Studies*, 24(1):1–45, 2011.

Gregory R Duffee. Forecasting with the term structure: The role of no-arbitrage restrictions. Technical report, Working paper, 2011a.

Gregory R Duffee. Information in (and not in) the term structure. *The Review of Financial Studies*, 24(9):2895–2934, 2011b.

Darrell Duffie, Jun Pan, and Kenneth Singleton. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, 68(6):1343–1376, 2000.

Daniel Egloff, Markus Leippold, and Liuren Wu. The term structure of variance swap rates and optimal variance swap investments. *Journal of Financial and Quantitative Analysis*, 45(5):1279–1310, 2010.

Robert Engle and Bryan Kelly. Dynamic equicorrelation. *Journal of Business & Economic Statistics*, 30(2):212–228, 2012.

Bjørn Eraker. Do stock prices and volatility jump? reconciling evidence from spot and option prices. *The Journal of Finance*, 59(3):1367–1403, 2004.

Bjørn Eraker, Michael Johannes, and Nicholas Polson. The impact of jumps in volatility and returns. *The Journal of Finance*, 58(3):1269–1300, 2003.

Peter Exterkate, Dick Van Dijk, Christiaan Heij, and Patrick JF Groenen. Forecasting the yield curve in a data-rich environment using the factor-augmented Nelson–Siegel model. *Journal of Forecasting*, 32(3):193–214, 2013.

Eugene F Fama and Kenneth R French. Dividend yields and expected stock returns. *Journal of Financial Economics*, 22(1):3–25, 1988.

Eugene F Fama and Kenneth R French. The cross-section of expected stock returns. *the Journal of Finance*, 47(2):427–465, 1992.

Eugene F Fama and Kenneth R French. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56, 1993.

Eugene F Fama and Kenneth R French. Multifactor explanations of asset pricing anomalies. *The Journal of Finance*, 51(1):55–84, 1996.

Eugene F Fama and James D MacBeth. Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy*, 81(3):607–636, 1973.

Jianqing Fan, Yingying Li, and Ke Yu. Vast volatility matrix estimation using high-frequency data for portfolio selection. *Journal of the American Statistical Association*, 107(497):412–428, 2012.

Carlo A Favero, Massimiliano Marcellino, and Francesca Neglia. Principal components at work: The empirical analysis of monetary policy with large data sets. *Journal of Applied Econometrics*, 20(5):603–620, 2005.

Carlo A Favero, Linlin Niu, and Luca Sala. Term structure forecasting: No-arbitrage restrictions versus large information set. *Journal of Forecasting*, 31(2):124–156, 2012.

Peter Feldhütter, Christian Heyerdahl-Larsen, and Philipp Illeditsch. Risk premia and volatilities in a nonlinear term structure model. *Review of Finance*, 22(1): 337–380, 2016.

Damir Filipović, Elise Gourier, and Loriano Mancini. Quadratic variance swap models. *Journal of Financial Economics*, 119(1):44–68, 2016.

Mario Forni, Marc Hallin, Marco Lippi, and Lucrezia Reichlin. The generalized dynamic factor model: One-sided estimation and forecasting. *Journal of the American Statistical Association*, 100(471):830–840, 2005.

Andras Fulop, Junye Li, and Jun Yu. Self-exciting jumps, learning, and asset pricing implications. *The Review of Financial Studies*, 28(3):876–912, 2014.

Anthony Garratt and Shaun P Vahey. Uk real-time macro data characteristics. *The Economic Journal*, 116(509):F119–F135, 2006.

Christian Genest, Kilani Ghoudi, and L-P Rivest. A semiparametric estimation procedure of dependence parameters in multivariate families of distributions. *Biometrika*, 82(3):543–552, 1995.

Raffaella Giacomini and Barbara Rossi. Forecast comparisons in unstable environments. *Journal of Applied Econometrics*, 25(4):595–620, 2010.

Domenico Giannone, Lucrezia Reichlin, and Luca Sala. Monetary policy in real time. *NBER Macroeconomics Annual*, 19:161–200, 2004.

Domenico Giannone, Michele Lenza, and Lucrezia Reichlin. Money, credit, monetary policy and the business cycle in the euro area: What has changed since the crisis? 2019.

William Goetzmann, Jonathan Ingersoll, Matthew Spiegel, and Ivo Welch. Portfolio performance manipulation and manipulation-proof performance measures. *The Review of Financial Studies*, 20(5):1503–1546, 2007.

Robin Greenwood and Dimitri Vayanos. Bond supply and excess bond returns. *The Review of Financial Studies*, 27(3):663–713, 2014.

Jaehoon Hahn and Hangyong Lee. Yield spreads as alternative risk factors for size and book-to-market. *Journal of Financial and Quantitative Analysis*, 41(2): 245–269, 2006.

Peter Reinhard Hansen. A test for superior predictive ability. *Journal of Business & Economic Statistics*, 23(4):365–380, 2005.

Nikolaus Hautsch, Lada M Kyj, and Roel CA Oomen. A blocking and regularization approach to high-dimensional realized covariance estimation. *Journal of Applied Econometrics*, 27(4):625–645, 2012.

Yi Hong and Xing Jin. Understanding common risk factors in variance swap rates, market return predictability and variance swap investments when volatility can jump. 2018.

John C Hull and Alan D White. Valuing credit derivatives using an implied copula approach. *The Journal of Derivatives*, 14(2):8–28, 2006.

Jean Jacod, Viktor Todorov, et al. Do price and volatility jump together? *The Annals of Applied Probability*, 20(4):1425–1469, 2010.

Harry Joe. *Multivariate models and multivariate dependence concepts*. CRC Press, 1997.

Scott Joslin, Anh Le, and Kenneth J Singleton. Why gaussian macro-finance term structure models are (nearly) unconstrained factor-VARs. *Journal of Financial Economics*, 109(3):604–622, 2013.

Scott Joslin, Marcel Priebsch, and Kenneth J Singleton. Risk premiums in dynamic term structure models with unspanned macro risks. *The Journal of Finance*, 69 (3):1197–1233, 2014.

Markus Junker, Alex Szimayer, and Niklas Wagner. Nonlinear term structure dependence: Copula functions, empirics, and risk implications. *Journal of Banking & Finance*, 30(4):1171–1199, 2006.

Don H Kim and Kenneth J Singleton. Term structure models and the zero bound: An empirical investigation of Japanese yields. *Journal of Econometrics*, 170(1): 32–49, 2012.

Don H Kim and Jonathan H Wright. An arbitrage-free three-factor term structure model and the recent behavior of long-term yields and distant-horizon forward rates. 2005.

Ralph SJ Koijen, Hanno Lustig, and Stijn Van Nieuwerburgh. The cross-section and time series of stock and bond returns. *Journal of Monetary Economics*, 88: 50–69, 2017.

Sadanori Konishi and Genshiro Kitagawa. *Information criteria and statistical modeling*. Springer Science & Business Media, 2008.

Siem Jan Koopman, Max IP Mallee, and Michel Van der Wel. Analyzing the term structure of interest rates using the dynamic Nelson–Siegel model with time-varying parameters. *Journal of Business & Economic Statistics*, 28(3):329–343, 2010.

Leo Krippner. Measuring the stance of monetary policy in zero lower bound environments. *Economics Letters*, 118(1):135–138, 2013.

Leo Krippner. A theoretical foundation for the Nelson–Siegel class of yield curve models. *Journal of Applied Econometrics*, 30(1):97–118, 2015.

Pavel Krupskii and Harry Joe. Factor copula models for multivariate data. *Journal of Multivariate Analysis*, 120:85–101, 2013.

Jean-Paul Laurent and Jon Gregory. Basket default swaps, CDOs and factor copulas. *Journal of Risk*, 7(4):103–122, 2005.

Anh Le and Kenneth J Singleton. The structure of risks in equilibrium affine models of bond yields. *Unpublished working paper, University of North Carolina at Chapel Hill.*, 2013.

Martin Lettau and Jessica A Wachter. The term structures of equity and interest rates. *Journal of Financial Economics*, 101(1):90–113, 2011.

Canlin Li, Linlin Niu, and Gengming Zeng. A generalized arbitrage-free Nelson-Siegel term structure model with macroeconomic fundamentals. *Available at SSRN 1787262*, 2011.

David X Li. On default correlation: A copula function approach. *The Journal of Fixed Income*, 9(4):43–54, 2000.

Junye Li and Gabriele Zinna. The variance risk premium: Components, term structures, and stock return predictability. *Journal of Business & Economic Statistics*, 36(3):411–425, 2018.

Yueh-Neng Lin. Pricing VIX futures: Evidence from integrated physical and risk-neutral probability measures. *Journal of Futures Markets: Futures, Options, and Other Derivative Products*, 27(12):1175–1217, 2007.

Rubén Loaiza-Maya and Michael Stanley Smith. Real-time macroeconomic forecasting with a heteroscedastic inversion copula. *Journal of Business & Economic Statistics*, pages 1–28, 2019.

Juan M Londono. The variance risk premium around the world. *Available at SSRN 2517020*, 2015.

André Lucas, Bernd Schwaab, and Xin Zhang. Conditional euro area sovereign default risk. *Journal of Business & Economic Statistics*, 32(2):271–284, 2014.

Sydney C Ludvigson and Serena Ng. Macro factors in bond risk premia. *The Review of Financial Studies*, 22(12):5027–5067, 2009.

Michael W McCracken and Joseph T McGillicuddy. An empirical investigation of direct and iterated multistep conditional forecasts. *Journal of Applied Econometrics*, 34(2):181–204, 2019.

Michael W McCracken and Serena Ng. FRED-MD: A monthly database for macroeconomic research. *Journal of Business & Economic Statistics*, 34(4):574–589, 2016.

Alexander J McNeil, Rüdiger Frey, Paul Embrechts, et al. *Quantitative risk management: Concepts, techniques and tools*, volume 3. Princeton university press Princeton, 2005.

Aleksey Min and Claudia Czado. Bayesian inference for multivariate copulas using pair-copula constructions. *Journal of Financial Econometrics*, 8(4):511–546, 2010.

Emanuel Mönch. Forecasting the yield curve in a data-rich environment: A no-arbitrage factor-augmented VAR approach. *Journal of Econometrics*, 146(1):26–43, 2008.

Roger B Nelsen. *An introduction to copulas*. Springer Science & Business Media, 2007.

Charles R Nelson and Andrew F Siegel. Parsimonious modeling of yield curves. *Journal of Business*, pages 473–489, 1987.

Dong Hwan Oh and Andrew J Patton. Modeling dependence in high dimensions with factor copulas. *Journal of Business & Economic Statistics*, 35(1):139–154, 2017.

Dong Hwan Oh and Andrew J Patton. Time-varying systemic risk: Evidence from a dynamic copula model of CDS spreads. *Journal of Business & Economic Statistics*, 36(2):181–195, 2018.

Alex D Patelis. Stock return predictability and the role of monetary policy. *the Journal of Finance*, 52(5):1951–1972, 1997.

Andrew Patton. Copula methods for forecasting multivariate time series. In *Handbook of Economic Forecasting*, volume 2, pages 899–960. Elsevier, 2013.

Andrew J Patton. A review of copula models for economic time series. *Journal of Multivariate Analysis*, 110:4–18, 2012.

H Hashem Pesaran and Yongcheol Shin. Generalized impulse response analysis in linear multivariate models. *Economics letters*, 58(1):17–29, 1998.

Ralitsa Petkova. Do the Fama–French factors proxy for innovations in predictive variables? *The Journal of Finance*, 61(2):581–612, 2006.

Davide Pettenuzzo, Allan Timmermann, and Rossen Valkanov. Forecasting stock returns under economic constraints. *Journal of Financial Economics*, 114(3): 517–553, 2014.

Marcel A Priebsch et al. A shadow rate model of intermediate-term policy rate expectations. Technical report, Board of Governors of the Federal Reserve System (US), 2017.

Patrick Royston. Approximating the Shapiro-Wilk W-test for non-normality. *Statistics and Computing*, 2(3):117–119, 1992.

Patrick Royston. A toolkit for testing for non-normality in complete and censored samples. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 42 (1):37–43, 1993.

Patrick Royston. Remark AS R94: A remark on algorithm AS 181: The W-test for normality. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 44(4):547–551, 1995.

Lucio Sarno, Paul Schneider, and Christian Wagner. The economic value of predicting bond risk premia. *Journal of Empirical Finance*, 37:247–267, 2016.

Samuel S Shapiro and RS Francia. An approximate analysis of variance test for normality. *Journal of the American Statistical Association*, 67(337):215–216, 1972.

Samuel Sanford Shapiro and Martin B Wilk. An analysis of variance test for normality (complete samples). *Biometrika*, 52(3/4):591–611, 1965.

Hideaki Shimazaki and Shigeru Shinomoto. Kernel bandwidth optimization in spike rate estimation. *Journal of Computational Neuroscience*, 29(1-2):171–182, 2010.

Galit Shmueli et al. To explain or to predict? *Statistical Science*, 25(3):289–310, 2010.

M Sklar. Fonctions de repartition an dimensions et leurs marges. *Publ. inst. statist. univ. Paris*, 8:229–231, 1959.

Michael S Smith and Shaun P Vahey. Asymmetric forecast densities for US macroeconomic variables from a Gaussian copula model of cross-sectional and serial dependence. *Journal of Business & Economic Statistics*, 34(3):416–434, 2016.

Michael Stanley Smith. Copula modelling of dependence in multivariate time series. *International Journal of Forecasting*, 31(3):815–833, 2015.

Michael Stanley Smith and Worapree Maneesoonthorn. Inversion copulas from nonlinear state space models with an application to inflation forecasting. *International Journal of Forecasting*, 34(3):389–407, 2018.

Elliott Sober. Instrumentalism, parsimony, and the Akaike framework. *Philosophy of Science*, 69(S3):S112–S123, 2002.

James H Stock and Mark W Watson. Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association*, 97 (460):1167–1179, 2002.

James H Stock and Mark W Watson. Disentangling the channels of the 2007–2009 recession. Technical report, National Bureau of Economic Research, 2012.

James H Stock and Mark W Watson. Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics. In *Handbook of Macroeconomics*, volume 2, pages 415–525. Elsevier, 2016.

Timo Teräsvirta. Forecasting economic variables with nonlinear models. *Handbook of Economic Forecasting*, 1:413–457, 2006.

Daniel L Thornton and Giorgio Valente. Out-of-sample predictions of bond excess returns and forward rates: An asset allocation perspective. *The Review of Financial Studies*, 25(10):3141–3168, 2012.

Hideatsu Tsukahara. Semiparametric estimation in copula models. *Canadian Journal of Statistics*, 33(3):357–375, 2005.

Dick Van Dijk, Siem Jan Koopman, Michel Van der Wel, and Jonathan H Wright. Forecasting interest rates with shifting endpoints. *Journal of Applied Econometrics*, 29(5):693–712, 2014.

Zhiguang Wang and Robert T Daigler. The performance of VIX option pricing models: Empirical evidence beyond simulation. *Journal of Futures Markets*, 31 (3):251–281, 2011.

Ivo Welch and Amit Goyal. A comprehensive look at the empirical performance of equity premium prediction. *The Review of Financial Studies*, 21(4):1455–1508, 2007.

Kenneth D West. Asymptotic inference about predictive ability. *Econometrica: Journal of the Econometric Society*, pages 1067–1084, 1996.

Halbert White. A reality check for data snooping. *Econometrica*, 68(5):1097–1126, 2000.

Jonathan H Wright. Term premia and inflation uncertainty: Empirical evidence from an international panel dataset. *American Economic Review*, 101(4):1514–34, 2011.

Jing Cynthia Wu and Fan Dora Xia. Measuring the macroeconomic impact of monetary policy at the zero lower bound. *Journal of Money, Credit and Banking*, 48(2-3):253–291, 2016.

Lilian Shiao-Yen Wu, Jeffrey S Pai, and JRM Hosking. An algorithm for estimating parameters of state-space models. *Statistics & Probability Letters*, 28(2):99–106, 1996.

Yuhang Xing, Xiaoyan Zhang, and Rui Zhao. What does the individual option volatility smirk tell us about future equity returns? *Journal of Financial and Quantitative Analysis*, 45(3):641–662, 2010.

Jin E Zhang and Yingzi Zhu. VIX futures. *Journal of Futures Markets: Futures, Options, and Other Derivative Products*, 26(6):521–531, 2006.

Yingzi Zhu and Jin E Zhang. Variance term structure and VIX futures pricing. *International Journal of Theoretical and Applied Finance*, 10(01):111–127, 2007.