Long Memory Conditional Volatility and Dynamic Asset Allocation

Submitted by Anh Thi Hoang Nguyen to the University of Exeter as a thesis for the degree of Doctor of Philosophy in Finance in September 2011.

I certify that all material in this thesis which is not my own work has been identified and that no material has previously been submitted and approved for the award of a degree by this or any other University.

(Signature) ………………………………………………………………………………

Acknowledgement

I would like to express my deep and sincere gratitude to my supervisor, Professor Richard Harris for his motivation and continuous support of my PhD research. His guidance helped me through all the time of research and writing of this thesis. I could not have imagined having a better supervisor and mentor for my PhD study.

I would also like to thank the Board of Examiners for their constructive comments and suggestions in my viva.

I am grateful to Gilles Zumbach for his useful comments on the analysis of the forecast performance of multivariate long memory volatility models.

I would like to thank the staff at the University of Exeter Business School, the IT officers and the librarians for their help in many different ways throughout my research.

I gratefully acknowledge the funding from the Ministry of Education and Training, Vietnam, that made my PhD work possible.

My time at Exeter was made enjoyable in large part due to many friends and groups that have become a part of my life. Thank you.

Lastly, and most importantly, I would like to thank my family for all their love, understanding and encouragement. To my parents, who raised me with a love of science and have been a constant source of support, and to my husband, who always stands by me and encourages me with his love, I dedicate this thesis.

Abstract

The thesis first evaluates the forecast performance of multivariate long memory conditional volatility models among themselves and against that of short memory conditional volatility models, using the asset allocation framework of Engle and Colacito (2006). While many alternative conditional volatility models have been developed in the literature, my choice reflects the need for parsimonious models that can be used to forecast high dimensional covariance matrices. In particular, I compare the statistical and economic performance of four multivariate long memory volatility models (the long memory EWMA, long memory EWMA-DCC, FIGARCH-DCC and Component GARCH-DCC models) with that of two multivariate short memory volatility models (the short memory EWMA and GARCH-DCC models). The research reports two main findings. First, for longer horizon forecasts, long memory volatility models generally produce forecasts of the covariance matrix that are statistically more accurate and informative, and economically more useful than those produced by short memory volatility models. Second, the two parsimonious long memory EWMA models outperform the other models – both short memory and long memory – in a majority of cases across all forecast horizons. These results apply to both low and high dimensional covariance matrices with both low and high correlation assets, and are robust to the choice of estimation window.

The multivariate conditional volatility models are then analysed further to shed light on the benefits of allowing for long memory volatility dynamics in forecasts of the covariance matrix for dynamic asset allocation. Specifically, the research evaluates the economic gains accruing to long memory volatility timing strategies, using the procedure of Fleming et al. (2001). The research consistently identifies the gains from incorporating long memory volatility dynamics in investment decisions. Investors are willing to pay to switch from the static to the dynamic strategies, and especially from the short memory volatility timing to the long memory volatility timing strategies across both short and long investment horizons. Among the long memory conditional volatility models, the two parsimonious long memory EWMA models, again, generally produce the most superior portfolios. When transaction costs are taken into account, the gains from the daily rebalanced dynamic portfolios deteriorate; however, it is still worth implementing the dynamic strategies at lower rebalancing frequencies. The results are robust to estimation error in expected returns, the choice of risk aversion coefficients and the use of a long-only constraint.

The long memory conditional covariance matrix is inevitably subject to estimation error. The research then employs a factor structure to control for estimation error in forecasts of the high dimensional covariance matrix. Specifically, the research develops a dynamic long memory factor (the Orthogonal Factor Long Memory, or OFLM) model by embedding the univariate long memory EWMA model of Zumbach (2006) into an orthogonal factor structure. The new factor model follows richer processes than normally assumed, in which both the factors and idiosyncratic shocks are modelled with long memory behaviour in their volatilities. The factor-structured OFLM model is evaluated against the six above multivariate conditional volatility models, especially the fully estimated multivariate long memory EWMA model of Zumbach (2009b), in terms of forecast performance and economic benefits. The results suggest that the OFLM model generally produces impressive forecasts over both short and long forecast horizons. In the volatility timing framework, portfolios constructed with the OFLM model consistently dominate the static and other dynamic volatility timing portfolios in all rebalancing frequencies. Particularly, the outperformance of the factor-structured OFLM model to the fully estimated LM-EWMA model confirms the advantage of the factor structure in reducing estimation error. The factor structure also significantly reduces transaction costs, making the dynamic strategies more feasible in practice. The dynamic factor long memory volatility model also consistently produces more superior portfolios than those produced by the traditional unconditional factor and the dynamic factor short memory volatility models.

Tables of Content

List of Tables

List of Figures

List of Appendices

List of Abbreviations

APT	Arbitrage Pricing Theory
AR	Autoregressive
ARCH	Autoregressive Conditional Heteroskedasticity
ARFIMA	Autoregressive Fractionally Integrated Moving Average
ARIMA	Autoregressive Integrated Moving Average
ARMA	Autoregressive Moving Average
BEKK	Baba, Engle, Kraft and Kroner
BIRR	Burmeister, Ibbotson, Roll and Ross
CAPM	Capital Asset Pricing Model
CCC	Constant Conditional Correlation
CGARCH	Component Generalised Autoregressive Conditional Heteroskedasticity
CML	Capital Market Line
CRR	Chen, Roll and Ross
DCC	Dynamic Conditional Correlation
DJIA	Dow Jones Industrial Averages
EGARCH	Exponential Generalised Autoregressive Conditional Heteroskedasticity
EWMA	Exponentially Weighted Moving Average
FIGARCH	Fractionally Integrated Generalised Autoregressive Conditional Heteroskedasticity
GARCH	Generalised Autoregressive Conditional Heteroskedasticity
GJR-GARCH	Glosten-Jagannathan-Runkle Generalised Autoregressive Conditional Heteroskedasticity
GMM	Generalised Method of Moments
GMV	Global Minimum Variance
GPH	Geweke-Porter-Hudak log periodgram estimator
HAC	Heteroscedasiticity and Autocorrelation Consistent
HML	High Minus Low
HMSE	Heteroscedasticity-adjusted Mean Squared Error
HYGARCH	Hyperbolic Generalised Autoregressive Conditional Heteroskedasticity
i.i.d	independently identically distributed

IC	Information Criterion
IGARCH	Integrated Generalised Autoregressive Conditional Heteroskedasticity
LM-EWMA	Long Memory Exponentially Weighted Moving Average
MAE	Mean Absolute Error
MS	Moulines-Soulier log periodgram estimator
MSE	Mean Squared Error
OFLM	Orthogonal Factor Long Memory
PCA	Principal Components Analysis
RMSE	Root Mean Squared Error
S&P500	Standard & Poor's 500 Index
SMB	Small Minus Big
TGARCH	Threshold Generalised Autoregressive Conditional Heteroskedasticity

Author's Declaration

I hereby declare that this thesis incorporates materials that are results of joint research, as follows:

Chapter 5 is based on a paper submitted to the International Journal of Forecasting co-authored with Professor Richard Harris. Professor Richard Harris provided editorial advice and guidance throughout the development of the analysis and the paper. Anh Nguyen carried out the analysis and wrote most of the paper.

Parts of Chapters 6 and 7 are based on a working paper co-authored with Professor Richard Harris. Professor Richard Harris provided editorial advice and guidance throughout the development of the model and the paper. Anh Nguyen developed the model, carried out the analysis and wrote most of the paper.

I am aware of the University of Exeter's regulation and I certify that I have properly acknowledged the contribution of other researchers to my thesis, and have obtained permission from them to include the above materials in my thesis.

I certify that, with the above qualification, this thesis, and the research to which it refers, is the product of my own work.

Chapter 1

Introduction

1.1 Background and Rationale

The variance-covariance matrix (hereinafter, the covariance matrix) plays a central role in many areas of applied finance, especially in asset allocation and risk management. In particular, the covariance matrix is one of the two inputs, along with expected returns, in the mean-variance asset allocation theory of Markowitz (1952, 1959). By incorporating risk and correlation in the asset allocation decision, Markowitz shifted the focus of the financial industry from single asset selection towards the concept of diversification through portfolio choice. While the expected return of a portfolio is the weighted average of the asset components' expected returns, its risk is always less than the weighted average risk of its components, unless the assets are perfectly positively correlated. Assets' risk and their interactions in a portfolio, implied through the notion of volatility and correlation, are hence vital in any investment decision.

The classic Markowitz's asset allocation theory assumes expected returns and the covariance matrix are known with certainty. However, these parameters are not observed in practice, and hence must be estimated. The traditional approach employs the sample covariance matrix as a proxy for the unknown true covariance matrix. The sample covariance matrix, though being the best unbiased and efficient estimator under the assumption of independently, identically distributed returns, is inherently subject to estimation error, especially when returns deviate from normality, which, unfortunately, is a prevailing feature in financial markets. Estimation error is multiplied in the high dimensionality of the covariance matrix typically used in asset allocation. The inversion of the covariance matrix may further aggravate the estimation error problem. Indeed, there is ample empirical evidence that shows the poor performance of optimal portfolios constructed with sample estimates (see, for example, Best and Grauer, 1991, Broadie, 1993, Britten-Jones, 1999). Estimation error of the sample estimates can be so serious that Michaud (1989) even calls the mean-variance optimiser "the estimation error maximiser." Equally-weighted portfolios are found to dominate mean-variance sample-based optimal portfolios in many cases (see, e.g., DeMiguel et al., 2009). Extensive research has thus been done to improve estimates and forecasts of the covariance matrix for robust asset allocation. One popular approach is to impose some structure on the covariance matrix, such as in the factor model or in the Bayesian-inspired shrinkage model (see Chan et al., 1999, Jagannathan and Ma, 2000, Ledoit and Wolf, 2003, 2004). Among others, Briner and Conner (2008) show in both simulation and empirical studies that a structured, though biased covariance matrix estimate may have better explanatory power than its unbiased, unstructured sample counterpart. Another promising direction is to exploit the high persistence of the time-varying conditional covariance matrix to generate better input forecasts for asset allocation. Andersen et al. (2006, p.794) acknowledge that "even if we rule out exploitable conditional mean dynamics, the optimal portfolio weights would still be time-varying from the second-order dynamics alone." Recent evidence suggests that dynamic asset allocation strategies, based on a time-varying conditional covariance matrix, dominate static strategies, based on the unconditional alternative (see, for example, Fleming et al., 2001, Han, 2006, Engle and Colacito, 2006, Thorp and Milunovich, 2007). Exploiting the predictability of the covariance matrix is now a key driver in asset allocation.

Dynamic asset allocation typically employs forecasts of the covariance matrix generated from popular conditional volatility models such as the multivariate Riskmetrics EWMA, multivariate GARCH, or multivariate Stochastic Volatility models, in which elements of the conditional covariance matrix are specified as weighted averages of the squares and cross-products of past return innovations with geometrically declining weights, so that shocks to variances and covariances dissipate rapidly. Consequently, dynamic strategies generally focus on short horizon investors who rebalance their portfolios daily. While this approach may make the most use of the forecast power of these conditional volatility models, it may not nevertheless correspond to the needs of investors, who often rebalance their portfolios at lower frequencies. Moreover, a mounting body of empirical evidence suggests that the autocorrelation functions of the squares and cross-products of returns decline more slowly than the geometric decay rate of the EWMA, GARCH and Stochastic Volatility models, and hence volatility shocks are more persistent than these models imply (see, for example, Taylor, 1986, Ding et al., 1993, Andersen et al., 2001). This 'long memory' feature is important not only for the measurement of current volatility, but also for forecasts of future volatility, especially over longer horizons. In particular, in the GARCH and Stochastic Volatility frameworks, forecasts of future volatility converge to the unconditional volatility at an exponential rate as the forecast horizon increases. In the EWMA framework, in contrast, a volatility shock has a permanent effect on forecast volatility at all horizons, and so forecasts of future volatility do not converge at all despite the fact that it is a short memory model. If volatility is indeed a long memory process, as the empirical evidence suggests, the short memory EWMA, GARCH and Stochastic Volatility models are misspecified. Moreover, the errors in forecasting the elements of the covariance matrix that arise from this misspecification are compounded as the forecast horizon increases.

The empirical evidence on volatility dynamics has prompted the development of long memory models of conditional volatility, and in the univariate context a number of approaches have been proposed. The FIGARCH model of Baillie et al. (1996) introduces long memory through a fractional difference operator, which gives rise to a slow hyperbolic decay for the weights on lagged squared return innovations while still yielding a strictly stationary process. The Hyperbolic GARCH (HYGARCH) model of Davidson (2004) is a generalisation that nests the GARCH, FIGARCH and IGARCH (or EWMA) models, allowing for a more flexible dynamic structure than the FIGARCH model and facilitating tests of short versus long memory in volatility dynamics. The Stochastic Volatility framework has been extended to allow for long memory by Breidt et al. (1998), who incorporate an ARFIMA process in the standard discrete time Stochastic Volatility model. Long memory can also be induced using a component structure for volatility dynamics. For example, the Component GARCH (CGARCH) model of Engle and Lee (1999) assumes that volatility is the sum of a highly persistent long run component and a mean-reverting short run component, each of which follows a short memory GARCH process. Similarly, Zumbach (2006) introduces a long memory model in which the dynamic process for volatility is defined as the logarithmically weighted sum of standard EWMA processes at different geometric time horizons. Like the short memory EWMA model of JP Morgan (1994) on which it is based, the long memory EWMA model has a highly parsimonious specification, which facilitates its implementation in practice.

In the multivariate context, long memory volatility modelling poses significant computational challenges, especially so for the high dimensional covariance matrices that are typically encountered in asset allocation and risk management. Indeed, so far the literature on long memory multivariate volatility modelling has restricted itself to the analysis of low dimensional covariance matrices, and has provided only limited evidence on the relative benefits from allowing for long memory in the multivariate setting. For example, Teyssiere (1998) estimates the covariance matrix for three foreign exchange return series using both an unrestricted multivariate FIGARCH model and a FIGARCH model implemented with the Constant Conditional Correlation (CCC) structure of Bollerslev (1990). Similarly, Niguez and Rubia (2006) model the covariance matrix of five foreign exchange series using an Orthogonal HYGARCH model, which combines the univariate HYGARCH long memory volatility model of Davidson (2004) with the multivariate Orthogonal GARCH framework of Alexander (2001). They show that the Orthogonal HYGARCH model outperforms the standard Orthogonal GARCH model in terms of one-day forecasts of the covariance matrix. Zumbach (2009b) develops a multivariate version of the univariate long memory EWMA model, in which elements of the covariance matrix are estimated as the averages of the squares and cross products of past returns with predetermined logarithmically decaying weights. This model is highly parsimonious and capable of handling large systems.

Allowing for long memory volatility dynamics in forecasts of the covariance matrix of returns may bring potential benefits for asset allocation over both short and long horizons. However, no studies have been done to explore the economic values of the long memory conditional covariance matrix for asset allocation. The research aims at filling this gap, studying the benefits of incorporating the long memory conditional covariance matrix in the asset allocation framework. Presumably, with slowly decaying autocorrelations, multivariate long memory volatility models are able to better capture the high persistence feature of volatility and covariance, and consequently exploit this feature to provide more reasonable forecasts of the covariance matrix over long horizons, which will potentially correspond more to the needs of most practical investors.

1.2 Research Questions and Scope

The overall aims of this research are twofold: (i) to evaluate the forecast performance of multivariate long memory conditional volatility models, and (ii) to examine the economic benefits that arise from allowing for long memory volatility dynamics in forecasting the covariance matrix in the asset allocation framework. Specifically, the research addresses the following questions:

1. Do multivariate long memory conditional volatility models produce better forecasts of the covariance matrix of returns than multivariate short memory conditional volatility models, especially for long horizons?
2. Are there any economic benefits when exploiting the long memory conditional covariance matrix for asset allocation, relative to using the unconditional or the short memory conditional alternatives?
3. How to control for estimation error in forecasts of the high dimensional long memory conditional covariance matrix?

The research is restricted to the analysis of multivariate conditional volatility models and their application to asset allocation. In particular, the research focuses on multivariate EWMA and multivariate GARCH models. Stochastic Volatility, Realised Volatility, Option Implied Volatility models are excluded from the research. Investors construct portfolios based on Markowitz's mean-variance asset allocation framework. For simplicity, I concentrate primarily on the single-period portfolio choice and ignore the hedging demands caused by time-varying investment opportunities. Portfolios comprise highly liquid assets, whose price data can be easily obtained at daily frequencies. These restrictions can, of course, be relaxed. However, this is beyond the scope of the thesis and is reserved for future research.

1.3 Contribution of the Thesis

The research evaluates the forecast performance of multivariate long memory conditional volatility models among themselves and against that of multivariate short memory conditional volatility models. While there exist a large number of conditional volatility models in the literature, my choice reflects the need for parsimonious models that can be used to forecast high dimensional covariance matrices. I employ four long memory volatility models: the multivariate long memory EWMA model of Zumbach (2009b), and three multivariate long memory volatility models implemented using the Dynamic Conditional Correlation (DCC) framework of Engle (2002). These are the univariate long memory EWMA model of Zumbach (2006), the Component GARCH model of Engle and Lee (1999) and the FIGARCH model of Baillie et al. (1996). This is the first study to evaluate the forecast performance of a range of multivariate long memory volatility models in high dimensional systems. The four multivariate long memory volatility models are compared with two multivariate short memory volatility models. These are the very widely used RiskMetrics EWMA model of JP Morgan

(1994), and the GARCH-DCC model. The six models are evaluated on the basis of both statistical and economic measures. While the statistical criteria examine the accuracy, bias and information content of the models' forecasts with measures such as the RMSE, MAE or Mincer-Zarnowitz regression, the economic criteria employ the economic loss function in the asset allocation framework of Engle and Colacito (2006). The research reports two main findings. First, for longer horizon forecasts, long memory volatility models generally produce forecasts of the covariance matrix that are statistically more accurate and informative, and economically more useful than those produced by short memory volatility models. Second, the two parsimonious long memory EWMA models outperform the other models – both short memory and long memory – in a majority of cases across all forecast horizons. These results apply to both low and high dimensional covariance matrices and both low and high correlation assets, and are robust to the choice of estimation window.

The multivariate conditional volatility models are then analysed further to shed light on the benefits of allowing for long memory volatility in estimating and forecasting the covariance matrix for dynamic asset allocation. Specifically, the research evaluates the economic value accruing to volatility timing strategies using the procedure of Fleming et al. (2001). The research consistently identifies the gains from incorporating long memory volatility dynamics in investment decisions. Investors are willing to pay to switch from the static to the dynamic volatility timing strategies, and from the short memory volatility to the long memory volatility models at both short and long investment horizons. Among the long memory conditional volatility models, the two parsimonious long memory EWMA models, again, generally produce the most superior portfolios. When transaction costs are taken into account, the gains from the daily rebalanced dynamic portfolios deteriorate; however, it is still worth implementing the dynamic strategies at lower rebalancing frequencies. The results are robust to estimation error in expected returns, the choice of risk aversion coefficient and the use of a long-only constraint.

The long memory conditional covariance matrix is inevitably subject to estimation error. The research then employs a factor structure to control for estimation error in forecasting the high dimensional long memory covariance matrix. In particular, the research develops a dynamic factor long memory conditional volatility (the Orthogonal Factor Long Memory, or OFLM) model by embedding the univariate long memory EWMA model of Zumbach (2006) into an orthogonal factor structure. The new factor model follows richer processes than normally assumed, in which both the factors and idiosyncratic shocks are modelled with long memory behaviour in their volatilities. The OFLM model is a generalisation of the Double Factor ARCH model of Engle (2009). The empirical results suggest that the dynamic factor OFLM model generally dominates the other multivariate conditional volatility models, both short memory and long memory, in terms of forecast performance and economic benefits across all forecast horizons. Especially, the outperformance of the factor-structured OFLM model to the fully estimated LM-EWMA model confirms the advantage of the factor structure in reducing estimation error. The factor structure also significantly reduces transaction costs, making the dynamic strategies more feasible in practice. The dynamic factor long memory volatility model also consistently produces more superior portfolios than those produced by the traditional unconditional factor and the dynamic factor short memory volatility models.

1.4 Structure of the Thesis

The thesis comprises eight chapters, beginning with this introductory chapter. Chapter 2 gives an overview of the classic asset allocation theory of Markowitz (1952, 1959). The chapter also provides a detailed analysis of the use of the unconditional covariance matrix for asset allocation.

Chapter 3 reviews time-varying conditional covariance matrix estimators and their application to asset allocation. The discussion especially focuses on the multivariate conditional volatility models that are applicable to a large number of assets. Advances in long memory conditional volatility models are highlighted.

Chapter 4 analyses the data. The research employs four sets of assets. These comprise a high correlation bivariate system (the S&P500 and DJIA indices), a low correlation bivariate system (the S&P500 and 10-year Treasury bond futures), and two moderate correlation high dimensional systems (the international stock and bond indices, and the components of the DJIA index).

Chapter 5 evaluates the forecast performance of long memory conditional volatility models. In particular, I compare the statistical and economic performance of four multivariate long memory volatility models (the long memory EWMA, long memory EWMA-DCC, FIGARCH-DCC and Component GARCH-DCC models) with that of two short memory volatility models (the short memory EWMA and GARCH-DCC models). The analysis covers investment horizons of up to 3 months and employs different estimation windows.

Chapter 6 examines the economic benefits of employing forecasts of the long memory conditional covariance matrix in asset allocation, using the volatility timing framework of Fleming et al. (2001). Portfolios constructed from the six multivariate conditional volatility models in Chapter 5 are evaluated using popular performance measures such as the out-of-sample Sharpe ratio, the abnormal return and the performance fee. Transaction costs are also taken into account. The robustness analysis studies the sensitivity of the findings to estimation error in expected returns, the choice of risk aversion coefficient and the use of a long-only constraint.

Chapter 7 turns to estimation error in forecasts of the high dimensional long memory conditional covariance matrix. A new Orthogonal Factor Long Memory conditional volatility model is developed by imposing a factor structure in the long memory volatility framework. The new factor-structured model is evaluated against the six multivariate conditional volatility models studied in the previous chapters, and especially against the fully estimated long memory EWMA model, in terms of both forecasting performance and economic benefits, using the procedures employed in Chapters 5 and 6. The performance of the Orthogonal Factor Long Memory volatility model is also compared with that of the traditional factor and the dynamic factor short memory volatility models.

The final chapter summarises the research, emphasizing all important findings. It also addresses the limitations of the research and suggests some implications for future studies.

Chapter 2

The Classical Asset Allocation Framework and Covariance Matrix Estimators

The seminal work of Markowitz (1952, 1959) laid the foundation for modern portfolio theory, providing a simple but powerful framework on how an optimizing investor would behave under uncertainty. Two fundamental economic insights, i.e., the concept of risk-return trade-off and the concept of diversification where risks are correlated, are beautifully captured in the classical asset allocation framework of Markowitz. While the benefits of diversification had been identified long before, Markowitz successfully translated the risk-return trade-off and diversification into an adequate theory of efficient portfolio investment, shifting the focus of the investment industry onto the interactions among securities in a portfolio. The classical asset allocation theory of Markowitz is reviewed in Section 2.1. Specifically, Section 2.1.1 introduces Markowitz's mean-variance portfolio optimisation theory, in which investors are concerned only with returns and risk in their portfolio choice decisions. The economic intuitions underlying the theory are illustrated graphically. Section 2.1.2 extends the basis framework, allowing investors to borrow and lend unlimitedly at a risk-free rate. In this case, investors can obtain better combinations than they can in the absence of the risk-free asset. Given the efficient frontiers built in the first two sections, Section 2.1.3 explains how investors choose their optimal portfolios. The decision is facilitated by the integration of utility theory into the model. Section 2.1.4 turns to the application of the mean-variance analysis in practice, where investors have to estimate expected returns and the covariance matrix for their investment decisions. Section 2.1.5 challenges the prohibitively restrictive assumptions of the classical model, introducing extensions in terms of, e.g., high-order moments of the return distribution, multiperiod investment horizons, and frictions such as transaction costs and taxation.

The estimation of the covariance matrix for asset allocation is discussed in detail in Section 2.2. The estimation of expected returns, which is beyond the scope of this research, can be found in, e.g., Jorion (1986), Fama and French (1992), Pesaran and

Timmermann (1995), Barberis (2000). In this chapter, the covariance matrix is assumed to be constant over time. The time-varying conditional covariance matrix will be studied in Chapter 3. Section 2.2.1 introduces the sample covariance matrix. Despite being the best unbiased and efficient estimator of the true covariance matrix, the sample covariance matrix is inevitably subject to estimation error, which is detrimental to the optimal portfolios. Mounting evidence of the poor out-of-sample performance of the optimal sample-based portfolios has prompted the development of improved covariance matrix estimators for robust asset allocation. A popular approach is to impose some structure on the covariance matrix. Section 2.2.2 details the factor structure. Section 2.2.2.1 sets the stage by introducing the linear factor decomposition. Single factor models are analysed in Section 2.2.2.2, whereas multifactor models are studied in Section 2.2.2.3. Section 2.2.2.4 investigates the practical implementation of factor models in the asset allocation problem, especially comparing the explanatory power among different types of factors. The Bayesian-inspired shrinkage structure is then described in section 2.2.3. Section 2.2.4 considers another structure, imposing a constant pairwise correlation coefficient onto the covariance matrix. Section 2.3 concludes the chapter.

The research in this chapter has indeed been well-documented. My contribution is simply to summarise existing knowledge on how to optimally allocate wealth in a portfolio. This chapter is, nevertheless, important for the thesis as it sets the stage for succeeding chapters.

2.1 The Classical Asset Allocation Framework

2.1.1 Markowitz's Mean-Variance Portfolio Optimisation Theory

Consider an investor who wants to construct an optimal portfolio at time t and hold it for a time horizon of Δt. Many assumptions lie behind the classical portfolio optimisation theory of Markowitz (1952, 1959). First, the investor is risk-averse; he requires a higher expected return to accept a higher level of risk. Therefore, we expect a positive relationship between expected returns and expected risk. In terms of portfolio choice, with a target expected return, the investor will choose the portfolio with the minimum risk from a set of all feasible portfolios. Markowitz is the first to quantify risk as the variance of the rate of return and use this risk measure to build his portfolio optimisation theory. Second, the investor is only concerned with returns and risk in his investment decision and constructs the optimal portfolio that is efficient in a return-risk, or mean-variance space. No higher-order moments of the return distribution are taken into account. Third, his investment horizon spans a single period Δt. He will not care about the gains or losses that may happen after the period Δt. He only updates his decision at time $t+\Delta t$. This behaviour is referred to as myopic (short-sighted) behaviour. For ease of notation, in the following I assume $\Delta t = 1$. Fourth, the assets are perfectly liquid; there are no transaction costs incurred when updating portfolios.

Denote $\mathbf{R}_t$ an *n*-dimensional vector of risky asset returns available for investment at time t and define $\boldsymbol{\mu}_{t+1} \equiv E_t\left(\mathbf{R}_{t+1}\right)$ as the vector of expected returns and $\Sigma_{t+1} = E_t[(\mathbf{R}_{t+1} - \boldsymbol{\mu}_{t+1})(\mathbf{R}_{t+1} - \boldsymbol{\mu}_{t+1})']$ as the covariance matrix of returns. Note the convention that the time subscripts are given for the date at which the variables are realised. Here, the investor is assumed to know $\boldsymbol{\mu}_{t+1}$ and Σ_{t+1} with certainty. His objective is to minimise risk with a target portfolio expected return $E_t(R_{p,t+1}) = \bar{\mu}$. He has to choose an $n \times 1$ vector of portfolio weights $\mathbf{w}_t$ to optimally allocate wealth among the *n* risky assets. Given the mean-variance optimisation framework, the investor solves the following constrained optimisation problem:

$$\min_{\mathbf{w}_t} \operatorname{var}(R_{p,t+1}) = \mathbf{w}_t' \Sigma_{t+1} \mathbf{w}_t \tag{2.1}$$

subject to

$$E_t(R_{p,t+1}) = \mathbf{w}_t'\boldsymbol{\mu}_{t+1} = \bar{\mu} \text{ and } \sum_{i=1}^{n} w_i = 1 . \tag{2.2}$$

The first constraint in (2.2) fixes the expected return of the portfolio to its target, while the second, a budget constraint, guarantees that all the wealth is invested. Short selling is permitted, which implies w_i can take negative values. Neither taxes nor transaction costs are included. Setting up the Lagrangian and the solution to the optimisation problem with equality constraints is easily obtained:

$$\mathbf{w}_t^* = \left(\frac{1}{AC - B^2}\right)\Sigma_{t+1}^{-1}(C\boldsymbol{\iota} - B\boldsymbol{\mu}_{t+1}) + \left(\frac{1}{AC - B^2}\right)\Sigma_t^{-1}(A\boldsymbol{\mu}_{t+1} - B\boldsymbol{\iota})\bar{\mu} , \tag{2.3}$$

where $\boldsymbol{\iota}$ denotes a unit vector, $A = \boldsymbol{\iota}'\Sigma_{t+1}^{-1}\boldsymbol{\iota}$, $B = \boldsymbol{\iota}'\Sigma_{t+1}^{-1}\boldsymbol{\mu}_{t+1}$, and $C = \boldsymbol{\mu}_{t+1}'\Sigma_{t+1}^{-1}\boldsymbol{\mu}_{t+1}$. The minimised portfolio variance is equal to $\mathbf{w}_t^{*'}\Sigma_{t+1}\mathbf{w}_t^*$.

If the investor follows the global minimum variance (GMV) strategy to construct the portfolio with the lowest possible risk regardless of the expected return, the optimal portfolio weights in (2.3) reduce to:

$$\mathbf{w}_t^* = \frac{1}{\boldsymbol{\iota}'\Sigma_{t+1}^{-1}\boldsymbol{\iota}}\Sigma_{t+1}^{-1}\boldsymbol{\iota}\,. \tag{2.4}$$

The optimisation problem described above is mathematically referred to as a quadratic optimisation problem. In this simple optimisation problem with equality constraints, solutions can be found analytically. However, in the more complex case when inequality constraints, e.g., non-negative portfolio weights, are imposed, numerical optimisation techniques may need to be applied to find the optimal solution.

We now look at a practical example. A rational investor wants to construct an efficient portfolio from 49 average-value-weighted industry portfolios of the US.[1] Results of the mean-variance efficient analysis are demonstrated in Figure 2.1. The vertical axis shows the annualised expected returns, while the horizontal axis shows the annualised risk as measured by returns' standard deviations. The blue envelope curve plots the efficient combinations of the 49 individual industry portfolios, which are marked with the purple stars. This set of combinations, which starts from the global minimum variance (GMV) portfolio (marked with the red circle) and going upwards, is called the *mean-variance efficient frontier*. The efficient frontier represents a set of portfolios that have the maximum expected return for a given level of risk, or the minimum level of risk for a target expected return $\left(\bar{\mu} > R_{GMV}\right)$. Portfolios located inside this frontier are dominated by others on the frontier that have the same risk but with higher expected return, or the same expected return but with lower risk, and so they are inefficient portfolios.

The mean-variance optimisation analysis implies two fundamental economic insights. First, it suggests the intuitive and powerful concept of diversification. Markowitz (1952, 1959) is the first to quantitatively evaluate assets not on their standalone performance basis, but on their interactions, through the notion of covariance, and contributions in a portfolio. While the expected return of a portfolio is the weighted average of the asset components' expected returns, its risk is always less than the weighted average risk of its components unless the assets are perfectly positively correlated. Figure 2.1 illustrates graphically the diversification effect. All the stars representing the individual industry

[1] Historical data of monthly returns of the 49 US industry portfolios are obtained from the Data Library of Kenneth R. French at http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

portfolios lie inside the blue curve, showing that the portfolios lie on the efficient frontier yield lower risk for a target expected return than the individual industry portfolios. Second, the framework captures the trade-off between expected return and risk. A portfolio is optimal if it has the highest expected return for a given level of risk or the lowest risk for a specified expected return. Starting with the least risky GMV portfolio at the left tip of the mean-variance efficient frontier, higher expected returns can only be achieved with higher levels of risk.

Alternative Formulations of the Mean-Variance Portfolio Optimisation Analysis

In the above section, we examine the investor who wants to minimise risk for a target expected return. Another perspective is to look at the decision that the investor has to make to maximise expected return when he cannot take more risk. This is one of the most commonly encountered problems in practice when, for example, portfolio managers are required to optimise with respect to tracking error volatility, i.e., the standard deviation of the difference between the portfolio's return and the benchmark return.

The investor now pre-determines a given level of risk $\bar{\sigma}^2$ and maximises his expected return of the portfolio. He solves the *maximum expected return* problem:

$$\max_{\mathbf{w}_t} E_t\left(R_{p,t+1}\right) = \mathbf{w}_t^{'}\boldsymbol{\mu}_{t+1} \tag{2.5}$$

$$\text{subject to the constraints of } \mathbf{w}_t^{'}\,\Sigma_{t+1}\,\mathbf{w}_t = \bar{\sigma}^2 \text{ and } \sum_i w_i = 1. \tag{2.6}$$

The mean-variance analysis can be formulated in an alternative way. Incorporating expected return and risk in a utility function in which the investor would prefer a high expected return with low variance portfolio, and the *maximum expected utility formulation* is given by

$$\max_{\mathbf{w}_t} E_t\left(U_{t+1}\right) = E\left(R_{p,t+1}\right) - \frac{\lambda}{2}\sigma_{p,t+1}^2 = \mathbf{w}_t^{'}\boldsymbol{\mu}_{t+1} - \frac{\lambda}{2}\mathbf{w}_t^{'}\,\Sigma_{t+1}\,\mathbf{w}_t \tag{2.7}$$

subject to $\sum_i w_i = 1$. Here, λ measures the investor's level of risk aversion. If diversification is agreed upon as the sound principle of investment, we would reject the objective of simply maximising expected return so that the investor would just invest all his wealth in one asset that generates the highest return. The risk aversion factor hence allows the investor to trade off mean and variance in a linear fashion. Any portfolio preferred by the investor will depend on his risk aversion. When the risk aversion coefficient λ is small, which means that the penalty for taking risk is small, the investor will choose more risky portfolios. In the extreme case when $\lambda = 0$, the return term dominates and the investor is willing to accept any level of risk in exchange for the highest return. Conversely, when the value of λ is large, risky portfolios will be highly penalised, leading to the choice of more conservative portfolios.

The three formulations as described above are equivalent in the sense that they all generate the same efficient frontier as they treat the trade-off between risk and return in a similar way, though from different standpoints.

2.1.2 Portfolio Choice Problem in the Presence of a Risk-Free Asset – The Capital Market Line

Suppose that there is a risk-free asset and the investor can borrow and lend unlimitedly at the risk-free rate R^f, he then can combine the optimal risky portfolio described above with the risk-free asset to create a superior portfolio. In the presence of the risk-free asset, the investor allocates a fraction $\mathbf{x}_t$ of his wealth to the *n* risky assets and the remainder $\left(1-\mathbf{x}_t^{'}\boldsymbol{\iota}\right)$ to the risk-free asset. Denote $\mathbf{r}_t \equiv \mathbf{R}_t - R^f$ the vector of excess returns over the risk-free asset. $\boldsymbol{\mu}_{t+1}$ and Σ_{t+1} are now defined as the expected returns and the covariance matrix of the excess returns, respectively. The expected portfolio return is then given by

$$E_t\left(R_{p,t+1}\right) = \mathbf{x}_t^{'}E_t\left(\mathbf{R}_{t+1}\right) + \left(1-\mathbf{x}_t^{'}\boldsymbol{\iota}\right)R^f = R^f + \mathbf{x}_t^{'}\left[E_t\left(\mathbf{R}_{t+1}\right) - R^f\right], \tag{2.8}$$

or in terms of excess returns,

$$E_t\left(R_{p,t+1}\right) = R^f + \mathbf{x}_t^{'}E_t\left(\mathbf{r}_{t+1}\right) = R^f + \mathbf{x}_t^{'}\boldsymbol{\mu}_{t+1}, \text{ or } E_t\left(r_{p,t+1}\right) = \mathbf{x}_t^{'}\boldsymbol{\mu}_{t+1}. \tag{2.9}$$

Given the mean-variance optimisation framework, the investor solves the following quadratic program, in terms of excess returns:

$$\min_{\mathbf{x}_t} \mathbf{x}_t^{'}\,\Sigma_{t+1}\,\mathbf{x}_t \tag{2.10}$$

$$\text{subject to } E_t\left(r_{p,t+1}\right) = \mathbf{x}_t^{'}\boldsymbol{\mu}_{t+1} = \overline{\mu}. \tag{2.11}$$

Note that no budget constraint is required. Since the investor can borrow or lend freely at the risk-free rate, the weights invested in the risky assets do not necessarily sum to one. The solution to this optimisation problem has a simpler form than in the case without a risk-free asset (see (2.3)):

$$\mathbf{x}_t^* = \frac{\bar{\mu} \Sigma_{t+1}^{-1} \boldsymbol{\mu}_{t+1}}{\boldsymbol{\mu}_{t+1}' \Sigma_{t+1}^{-1} \boldsymbol{\mu}_{t+1}} . \tag{2.12}$$

Other formulations can be represented in terms of excess returns in a similar way. For a particular choice such that $\sum_i x_i = 0$, all the wealth is invested in the risk-free asset. On the other hand, if $\sum_i x_i = 1$, the portfolio consists of all risky assets. Graphically, the mean-variance efficient frontier is now the straight line from the risk-free rate on the vertical axis and tangent to the old curved efficient frontier (Figure 2.2). The point where the straight line touches the curved line is called the *tangency* portfolio, which consists of all risky assets. This is the best mix of the risky assets that maximises the Sharpe ratio, defined as the sample mean of the realised portfolio excess returns over the risk-free rate divided by their sample standard deviation, and represented graphically by the slope of straight line. As the straight line lies above the curved line, the investor can obtain the target expected return with a lower level of risk than he can in the absence of the risk-free asset. The investor can even move aggressively further along the straight line past the tangency portfolio by borrowing at the risk-free rate to construct a leveraged portfolio. The straight line is popularly referred to as the Capital Market Line (CML). The discussion of the CML suggests that the investor optimally combines the risk-free asset with the same portfolio of risky assets - the tangency portfolio. Depending on his attitude towards risk, he can conservatively move down and to the left, or aggressively move up and to the right of the Capital Market Line. However, he should not alter the relative ratios of the risky assets in the tangency portfolio.

2.1.3 Incorporating Utility Theory – How to Invest Optimally?

Given the mean-variance efficient frontier or the CML (in the presence of a risk-free asset), how can the investor choose his optimal portfolio? This decision depends largely on his tolerance for risk. Each investor has different preferences and attitudes towards risk, thus choosing a different optimal portfolio. Utility functions can be incorporated into the framework to help understand how the investor allocates his funds optimally when faced with a set of possible choices.

Utility functions can be illustrated graphically in the form of indifference curves. Different indifference curves correspond to different levels of utility. Moving up from one indifference curve to another shows higher utility, while moving along one indifference curve just shows different combinations for the same level of utility. Figure 2.3 illustrates the portfolio choice decision when utilities (indifference curves) are included. The investor will choose the optimal portfolio at point A, which is the tangency of the indifference curve U_1 with the CML, or at point B, the tangency of the indifference curve U_2 with the mean-variance frontier. Note that, in the absence of a risk-free asset, the investor can only reach point B, obtaining a lower utility level than he can with the presence of a risk-free asset.

Utility functions allow us to generalise the mean-variance framework into a much wider class of problems, expected utility maximisation problems. Assume that an investor with utility defined over initial wealth $U(W_0)$ wants to maximise his expected utility with the end-period wealth $W = (1+R_p)W_0$. Also assume that there is no risk-free asset.

$$\max_{\mathbf{w}} EU(W) = E\left[U\left((1+R_p)W_0\right)\right] = W_0\left[1+E\left(U\left(R_p\right)\right)\right] \tag{2.13}$$

$$\text{subject to } \sum_i w_i = 1. \tag{2.14}$$

The budget constraint (2.14) is removed if there is a risk-free asset. Applying a Taylor series expansion of $U\left(R_p\right)$ around the mean $E\left(R_p\right) = \mu_p$, we have:

$$U\left(R_p\right) = U\left(\mu_p\right) + \left(R_p - \mu_p\right)U'\left(\mu_p\right) + \frac{1}{2}\left(R_p - \mu_p\right)^2 U''\left(\mu_p\right) + \text{higher-order terms} \tag{2.15}$$

Taking expectations on the both sides yields:

$$E\left[U\left(R_p\right)\right] = U\left(\mu_p\right) + \frac{1}{2}\sigma_p^2 U''\left(\mu_p\right) + E\left(\text{higher-order terms}\right) \tag{2.16}$$

as $E\left(R_p - \mu_p\right) = 0$ and $E\left(R_p - \mu_p\right)^2 = \sigma_p^2$.

It can be inferred from (2.16) that the expected utility maximisation is equivalent to the mean-variance optimisation in two special cases. First, asset returns are jointly elliptically distributed, i.e., fully described by the first two moments (Owen and Rabinovitch, 1983). In this case, $EU(W)=\left[1+E\left(U\left(R_p\right)\right)\right]W_0$ is just a function of portfolio mean and variance, no matter what utility function the investor may adopt. Second, the investor has a quadratic utility function, in which case the expectation of the higher-order terms vanishes. No distributional assumption is required on asset returns. For other utility functions, the mean-variance optimisation can only at best be interpreted as a second-order approximation of expected utility maximisation.

The investor's decision making process can now be divided into two separate stages. This is known as the Separation Theorem of Tobin (1958). First, the investor uses his knowledge about assets' expected returns and covariance matrix to derive the efficient frontier and the CML. This process is the same for all investors irrespective of their preferences. If expectations are assumed to be homogenous across all investors, then they should hold the same proportion of these risky assets. In the second stage, based on his subjective risk-return preference, the investor will choose which point on the CML to invest. If he is very risk averse, he will put most of his wealth in the risk-free asset and little in the risky assets. On the contrary, an investor with higher risk tolerance will invest more in the risky assets. In this stage, his subjective preference will not affect the relative weights among the risky assets derived in the first stage.

2.1.4 Application of the Classical Asset Allocation Theory

Markowitz optimisation treats expected returns, variances and covariances as deterministic. However, in practice, these moments of returns are unobservable and must be estimated. The sample moment estimates are typically employed as proxies for the unknown true parameters. Statistical estimates, nevertheless, are subject to estimation error. Define the estimated mean $\hat{\boldsymbol{\mu}}$ and covariance matrix $\hat{\Sigma}$ in a general way: $\hat{\boldsymbol{\mu}}=\boldsymbol{\mu}+\hat{\boldsymbol{\varepsilon}}_{\boldsymbol{\mu}}$ and $\hat{\Sigma}=\Sigma+\hat{\boldsymbol{\varepsilon}}_{\Sigma}$, where $\hat{\boldsymbol{\varepsilon}}_{\boldsymbol{\mu}}$ is an $n\times 1$ vector and $\hat{\boldsymbol{\varepsilon}}_{\Sigma}$ is an $n\times n$ matrix, both representing estimation error. Employing the expected sample proxies, the mean-variance quadratic optimisation becomes:

$$\begin{aligned}\max_{\mathbf{x}} E(U) &= \mathbf{x}'\hat{\boldsymbol{\mu}}-\frac{\lambda}{2}\mathbf{x}'\hat{\Sigma}\mathbf{x} \\ &= \mathbf{x}'\boldsymbol{\mu}+\mathbf{x}'\hat{\boldsymbol{\varepsilon}}_{\boldsymbol{\mu}}-\frac{\lambda}{2}[\mathbf{x}'\Sigma\mathbf{x}+\mathbf{x}'\hat{\boldsymbol{\varepsilon}}_{\Sigma}\mathbf{x}].\end{aligned} \tag{2.17}$$

It can be seen from (2.17) that the portfolio optimisation may be distorted. Instead of trading off the true portfolio expected return $\mathbf{w}'\boldsymbol{\mu}$ against its true variance $\mathbf{w}'\Sigma\mathbf{w}$, we trade off the true expected return plus the estimation error $\mathbf{w}'\hat{\boldsymbol{\varepsilon}}_{\boldsymbol{\mu}}$ against the true variance plus the estimation error $\mathbf{w}'\hat{\boldsymbol{\varepsilon}}_{\Sigma}\mathbf{w}$. Two sources of estimation error, in the mean vector and in the covariance matrix, both interact with the true values.

The expected optimal weights under (2.17) are:

$$E\left(\hat{\mathbf{x}}^*\right)=\frac{1}{\lambda}E\left(\hat{\Sigma}^{-1}\right)E\left(\hat{\boldsymbol{\mu}}\right)=\frac{1}{\lambda}E\left[\left(\Sigma+\hat{\boldsymbol{\varepsilon}}_{\Sigma}\right)^{-1}\right]\left[\boldsymbol{\mu}+E\left(\hat{\boldsymbol{\varepsilon}}_{\boldsymbol{\mu}}\right)\right]. \tag{2.18}$$

Under the i.i.d assumption, $E\left(\hat{\boldsymbol{\varepsilon}}_{\Sigma}\right)=0$ and $E\left(\hat{\boldsymbol{\varepsilon}}_{\boldsymbol{\mu}}\right)=0$, and so the optimal weights derived using the sample moments are the unbiased estimates of the true weights. However, most financial time series are not independently and identically distributed. As a result, non-trivial estimation error will be fed through to portfolio weights, distorting the optimisation results. Britten-Jones (1999) derives the exact finite sample inference procedure for hypothesis about the weights of efficient portfolios. Applying this procedure to an international mean-variance efficient portfolio, he finds excessive sampling error in the estimates of the optimal weights. The literature also shows that optimisation may produce extreme and non-intuitive weights for some assets, which contradicts the notion of diversification. Furthermore, the optimal solution may be highly unstable. The weights calculated from the mean-variance optimisation can be very sensitive to changes in the two input parameters, the expected return vector and the covariance matrix. As the mean-variance analysis exploits the smallest difference, small changes in the inputs can lead to dramatic changes in the portfolio weights (see Best and Grauer, 1991). Michaud (1989) even calls the mean-variance optimiser “the estimation error maximiser,” arguing that mean-variance optimisation may significantly overweight securities with large expected returns, low correlations and low variances, which are also the ones with the most estimation error.

To study the problem of estimation error in more detail, I follow Jobson and Korkie (1981) and Broadie’s (1993) experiments. Assume an investor wants to optimally allocate his wealth among 10 US industry portfolios.[2] I set the sample mean and covariance matrix to be the true parameters and simulate independent sets of 250 hypothetical monthly return samples of different sample sizes (60 months and 120

[2] Monthly returns on 10 US industry portfolios (Jan 1996 – Dec 2010) are obtained from the Data Library of Kenneth French at http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

months) from a multivariate normal distribution with the true moments. For each hypothetical set of returns, I compute the estimated parameters to construct the *estimated frontier*. The *true frontier* is constructed from the true parameters. Following Broadie, I also generate the *actual frontiers*, which are achieved by applying the portfolio weights derived from the estimated parameters to the true parameters. According to Broadie, the estimated frontier is what *appears* to be the case based on the data and estimated parameters, but the actual frontier is what *really* occurs based on the true parameters. The true frontier and the average estimated and actual frontiers are plotted in Figure 2.4. The results suggest that for any level of risk, the estimated frontier overestimates the expected return, hence shifting the average estimated curve upward and to the right relative to the true position. The distance between the two curves widens with higher levels of risk and expected returns (see, in particular, the dislocation of the upper right-hand point of the true frontier to the corresponding endpoint of the average estimated frontier). The estimated frontier tends to exaggerate certain error in the input parameters, resulting in optimistically biased estimates of portfolio performance (Broadie, 1993). In general, the estimated frontier tends to overestimate expected returns and underestimate risk of a portfolio. On the contrary, the actual frontier lies below the true frontier, and farther below the estimated frontier. The actual frontier can be interpreted as the out-of-sample frontier where the estimated parameters are employed to derive the portfolio weights, which are then applied to the realised true parameters to compute the realised portfolio performance. The difference in performance between the estimated and the actual frontiers illustrates the difference in the optimistic in-sample and the dismal out-of-sample performance of the mean-variance optimal portfolios in practice.

Figure 2.5 sheds more light on the poor performance of the actual frontiers. The blue solid curve is the true efficient frontier, while the red dashed curves are the 250 simulated actual frontiers. The actual frontiers are extremely volatile and consistently inferior to the true frontier. Increasing the sample size reduces the volatility of the actual frontiers, but cannot eliminate the problem.

Despite the simple and intuitive appeal of Markowitz's mean-variance optimisation, its application is often problematic. Extensive research has been done to provide resolutions to the well-documented practical problems associated with the mean-variance framework. Studies focus on reserving the benefits of the traditional framework while enhancing its practical value and effectiveness. Improved estimates of expected returns and the covariance matrix have been suggested, ranging from factor models to shrinkage estimators. Factor models are based on the evidence that asset returns have common risk factors that act as main sources of common variation among asset returns. This intuition is justified in some finance theories such as the Capital Asset Pricing or the Arbitrage Pricing Theory models. Those asset pricing models are useful not only in estimating asset expected returns but also in reducing the dimensionality of the covariance matrix. The idea of shrinkage estimators is pioneered by James and Stein (1961), who suggest an estimator that "shrinks" the sample mean toward a common "grand" mean across all variables. By shrinking the most extreme coefficients of the sample mean towards more central values, estimation error is suggested to be systematically reduced when it matters most. Mean shrinkage estimators have been applied to portfolio optimisation by Frost and Savarino (1986) and Jorion (1986), to name a few. For example, Jorion shows theoretically and in a simulation analysis that the James-Stein shrinkage estimator has lower estimation error than the sample mean. Also based on the idea of James and Stein, Ledoit and Wolf (2003, 2004) generalise the mean shrinkage estimator to the covariance matrix. Setting in a global minimum variance portfolio framework to abstract from the problem of estimating expected returns, they show that their shrinkage estimator produces portfolios with significantly lower out-of-sample variances than those produced by a set of well-established competing approaches.

Constraints on portfolio weights, such as no short sales or upper bounds, can be imposed on portfolios. These constraints are relevant in practice, though investment practitioners are normally faced with more constraints. Constraints are useful for controlling portfolio structure, hence reducing estimation error. Frost and Savarino (1988) demonstrate that portfolio constraints truncate extreme portfolio weights and thereby improve portfolio performance. Jagannathan and Ma (2000) go further when interpreting the constraints under certain conditions as a form of shrinkage estimation that improves the efficiency of the optimal portfolio. Their Monte Carlo simulations and empirical tests show that with nonnegative weight constraints in place, global minimum variance and minimum tracking error portfolios constructed using the sample covariance matrix perform as well as those constructed using factor and shrinkage models. However, Green and Hollifield (1992) suggest that extreme weighting is likely to be attributable to the dominance of a single factor in equity returns, which is equally true for population and estimated moments of returns. Thus, imposing weight constraints may produce specification error. Due to the trade-off between specification and estimation error, constraints may be only useful when estimation error is excessive.

The investor can follow an alternative approach, in which he chooses the portfolio weights that are optimal with respect to his subjective belief about the true return distribution. Given parameter uncertainty, he builds his subjective distribution of asset returns based on his prior belief about the true parameters and on the data he observes, using a Bayesian procedure. He then solves an average optimisation problem over all possible sets of parameter values derived from his subjective return distribution, where the expected utility of any given set of parameter values is weighted by his subjective probability of these parameter values. In this Bayesian approach, priors are of utmost importance. Priors can be uninformative, which contain little information about the parameters and lead to results that are comparable, but not identical in finite samples, to sample estimates (see, for example, Barry, 1974, Klein and Bawa, 1976, Brown, 1979). Later research specifies priors to rely on the theoretical implications of economic models. Priors can be the risk premia implied in the mean-variance optimisation theory and market equilibrium (Black and Litterman, 1992), the belief in market efficiency (Kandel and Stambaugh, 1996), or the belief in an asset pricing model (Pástor, 2000).

2.1.5 Relaxation of the Assumptions

The pioneering work of Markowitz provides a convenient and practical framework for asset allocation, based only on expected returns and the covariance matrix. The appealing simplicity of the model is achieved, however, with a set of prohibitively restrictive assumptions, which inherently hinders its application in practice. Considerable effort has thus been devoted to relaxing those assumptions. One line of research is to capture the preferences towards higher-order moments of returns and incorporating higher-order utility preferences in investors' objective functions. Other directions include, but are not restricted to, extensions to multiperiod investment horizons, and analysis of the effects of frictions in the investment decision-making process. This section provides a brief summary of some advances in the portfolio choice problem.

Portfolio Choice with Higher Order Moments

Mounting evidence claims that the problem of maximising investors' expected utilities cannot always be reduced to the problem of mean-variance trade-off (see, for example, Samuelson, 1970, Kraus and Litzenberger, 1976). The mean-variance analysis is only a special case of expected utility maximisation that arises when asset returns are elliptically distributed or when investors have a quadratic utility function. However, it is well established that financial return's distribution generally cannot be fully characterised by the mean and variance alone. Asset returns typically have fatter tails than those implied in the normal distribution and are often not symmetric (see Mandelbrot, 1963, Fama, 1965). The literature also suggests that the fat tails and skewness of returns may affect investors' decisions in allocating wealth; investors generally exhibit preference for positively skewed and light-tailed to negatively skewed and heavy-tailed asset return distributions. Incorporating conditional skewness in an asset pricing problem, Harvey and Siddique (2000) show that non-increasing absolute risk aversion, a critical feature of risk-averse individuals, implies a preference for high skewness. Adding an asset with negative coskewness will reduce total portfolio skewness, leading investors to require a higher expected return than that required when they add an asset with identical risk characteristics but with zero coskewness. Similarly, Lai (1991) and Chunhachinda et al. (1997) solve a multi-objective portfolio choice problem (i.e., maximising expected returns and skewness with a specified level of risk) and suggest that investors trade expected returns for skewness. Incorporating higher-order moment preference into the asset allocation framework may require the extension of utility functions. Studies typically apply the Taylor series expansion to derive higher-order approximations of expected utility functions (see, for example, Brandt et al., 2005, Jondeau and Rockinger, 2006). Uncertainty in parameters is also taken into account. Harvey et al. (2010) embed a multivariate skew normal model in a Bayesian framework to address the parameter uncertainty of higher-order moments. Similarly, concerned about the sensitivity of the conventional moments to outliers, Jurczenko et al. (2008) advocate the use of L-moments, deriving optimal portfolios in a four-dimensional non-convex mean-L-variance-L-skewness-L-kurtosis space and presenting various illustrations of the first four L-moment efficient portfolios. All these studies suggest the importance of integrating high-order moments into portfolio selection, especially when returns show strong deviation from the elliptical distribution.

Incorporating higher-order moments, on the one hand, allows the mean-variance framework to better reflect the characteristics of asset returns observed in practice. On the other hand, it makes the practical implementation much more complicated. Allowing for higher-order moments implies more parameters to be estimated. High dimensionality, which is already a serious concern in the context of covariance matrix estimation (more details are given on Section 2.2.1), is more problematic when coskewness and cokurtosis parameters are involved. For example, optimising a 10-asset portfolio requires the estimation of 55 variance-covariance, 220 skewness-coskewness, and 715 kurtosis-cokurtosis parameters! As a result, the research on high-order moments in asset allocation generally restricts itself to very low dimensional systems. High dimensionality also generates excessive estimation error. To reduce estimation error in estimating higher-order moments, Martellini and Ziemann (2010) extend to higher-order comovement several models that have been extensively applied to reduce dimensionality and estimation error in the covariance matrix. They find that portfolios with improved higher-order estimates yield superior performance to those with sample estimates. Improving estimates of the high-order moments promises an interesting direction of research. Given the remarkable increase in dimensionality, estimation error must be controlled; otherwise, they may be so large that they may offset all the gains from a more correctly specified framework.

Multiperiod Investment Horizons

The Markowitz process is a single-period portfolio choice, while real-world practice normally requires longer horizons with intermediate rebalancing. Extensive research has been done to formulate the portfolio choice problem as an intertemporal expected utility maximisation (see, for example, Samuelson, 1969, Merton, 1971, Merton, 1973). In both discrete and continuous time formulation, the literature shows that dynamic intertemporal optimal portfolio choice in a multiperiod context can be substantially different from a sequence of myopic single-period portfolio choices in terms of asset allocation and expected utilities (see Brandt, 2009, for a detailed analysis). The difference is termed the *hedging demands* as investors try to hedge against changes in investment opportunities when deviating from the single period portfolio choice. The classic results of Samuelson (1969) and Merton (1971) derive two restrictive conditions under which a long-term investor may act myopically, choosing the same portfolios as a short-term investor: (i) the investor has constant investment opportunities so that he does not need to hedge (an obvious case is power utility and independently, identically distributed returns), or (ii) the investor has log utility (in this case returns are not required to be i.i.d). However, investment opportunities are not constant as real interest rates move over time, and even if expected excess returns on risky assets over risk-free assets are constant, time variation in real interest rates is enough to generate large differences between optimal portfolios for long-term and short-term investors

(Campbell and Viceira, 2002). Short-term and long-term investors are also different in the sense that short-term investors' wealth is assumed to consist of only tradable financial assets, which is not realistically true for long-term individual investors. Long-term individual investors, who are working and saving for, among others, retirement, own tradable financial assets just as part of their total wealth. They also own a very valuable untradable asset, their labour income. The introduction of labour income in asset allocation has prompted numerous extensions of the theory (see, for example, Heaton and Lucas, 2000, Viceira, 2001, Campbell et al., 2001, Letendre and Smith, 2001, Cocco et al., 2005, Angerer and Lam, 2009) .

Transaction Costs and Taxation

The classic framework can be modified to allow for frictions such as transaction costs and taxation. Almost all portfolios require some adjustments during their lifetime, hence incurring non-trivial transaction costs. In a continuous time setting, Davis and Norman (1990) study a one-risky-asset portfolio where there are charges on all transactions equal to a fixed percentage of the amount traded and derive the exact algorithm to solve the optimal policies. In a multiple risky asset context, Leland (2000) studies a single-period investor who minimises the sum of the proportional transaction costs and the variance of the tracking error. He develops a numerical procedure to calculate the optimal rebalancing rule and implements the procedure in a number of examples. Also working with multiple risky assets but in a multiperiod problem with predictable returns, Lynch and Tan (2003) develop methods to numerically solve investors' decision making problem when transaction costs are accounted for. They also perform some utility comparisons, including the assessment of the utility cost of transaction costs.

Taxation is another friction faced by investors when making investment decisions. For example, selling assets generates capital gains tax. Incorporating tax in portfolio choice is extremely difficult in the context of the realisation-based feature of tax and of complex myriads of tax codes for different types of transactions and investors. The usual approach is to adopt the most significant features of the tax code and to assume the other unmodelled features are of secondary importance. Recent papers on the implications of taxation on portfolio choice include Dammon et al. (2001), DeMiguel and Uppal (2005), Garlappi and Huang (2006), Huang (2008).

2.2 Covariance Matrix Estimators

Markowitz's mean-variance portfolio analysis requires the estimation of expected returns and the covariance matrix. This section presents a detailed discussion of the estimation of the covariance matrix for asset allocation. The review is restricted to unconditional covariance matrix estimators. Time-varying conditional covariance matrix estimators will be investigated in Chapter 3. The estimation of expected returns, which is beyond the scope of this research, can be found in, e.g., Jorion (1986), Fama and French (1992), Pesaran and Timmermann (1995), Barberis (2000).

2.2.1 The Sample Covariance Matrix Estimator

Consider an $n \times T$ matrix of excess returns $\mathbf{r}$ on n assets over a sample of T periods. Denote $\boldsymbol{\mu}$ and Σ the mean and covariance matrix of the excess returns, respectively. The sample unconditional estimators of the mean $\hat{\boldsymbol{\mu}}$ and covariance matrix $\hat{\Sigma}$ are:

$$\hat{\boldsymbol{\mu}} = \frac{1}{T}\mathbf{r}\boldsymbol{\iota}', \tag{2.19}$$

$$\hat{\Sigma} = \frac{1}{T-1}\tilde{\mathbf{r}}\tilde{\mathbf{r}}', \tag{2.20}$$

where $\boldsymbol{\iota}$ is an $1 \times T$ unit vector, and $\tilde{\mathbf{r}}$ is the matrix of mean-corrected returns:

$$\tilde{\mathbf{r}} = \mathbf{r} - \hat{\boldsymbol{\mu}}\boldsymbol{\iota} = \mathbf{r} - \left(\frac{1}{T}\mathbf{r}\boldsymbol{\iota}'\right)\boldsymbol{\iota} = \mathbf{r}\left(\mathbf{I} - \frac{1}{T}\boldsymbol{\iota}'\boldsymbol{\iota}\right) \tag{2.21}$$

with $\mathbf{I}$ being a $T \times T$ unit matrix.

It is obvious that $rank\left(\hat{\Sigma}\right) = \min\left(n, T-1\right)$. Hence, for n assets, we need at least $n+1$ periods if we want the sample covariance matrix to be invertible, a requirement to estimate the weights of the efficient portfolio.

If returns are i.i.d, then the sample covariance matrix has the appealing property of being the best unbiased estimate of the true covariance matrix. However, the sample covariance matrix is inevitably error-prone. To get a sense of estimation error in the covariance matrix, I repeat the experiments in Section 2.1.4. Specifically, I consider an optimisation simulation, in which the investor knows the true expected returns but uses the sample estimates of the covariance matrix. Figure 2.6 shows the histograms of the

Sharpe ratios of the *actual* tangency portfolios for 24 and 120 monthly observations. The vertical line in each plot represents the Sharpe ratio of the *true* tangency portfolio. The Sharpe ratios of the simulated portfolios are volatile and lower than the true Sharpe ratio, especially in the case of $T = 24$ (recall that $n = 10$). When the sample size T increases relatively to the number of assets $(T = 120)$, estimation error reduces as there are more observations per parameter. The performance of the actual tangency portfolios is then improved.

In practice, we rarely enjoy the luxury of having the number of observations significantly larger than the number of assets $T >> n$. It is normal that a portfolio consists of hundreds of assets while the sample period is bounded by a few years. When T is not greatly larger than n, the sample covariance matrix may not be well-conditioned, yielding huge estimation error when being inverted. Lengthening the sample period is problematic since observations far in the distant past may have little explanatory power relative to current observations. Dimensionality is another problem. The number of estimated parameters increases with the square of the number of assets. For an n-asset portfolio, we have to estimate a covariance matrix of $\frac{1}{2}n(n+1)$ parameters. For instance, if we have 100 assets to choose from, we have to estimate 2025 parameters of the covariance matrix. The more parameters to be estimated, the more estimation error is likely to arise. Estimation error may be so excessive that it renders the optimal portfolio practically worthless and difficult to understand.

2.2.2 Factor Models

One popular approach to reduce estimation error in the covariance matrix is to impose a factor structure on the covariations among assets. The factor structure reduces the number of parameters to be estimated, and hence reduces estimation error. However, it comes at a price. The structured covariance matrix with a few factors may not capture every relationship among assets, incurring specification error. Increasing the number of factors reduces specification error, yet with an increase in estimation error. Selecting the ‘optimal’ factors involves a trade-off between estimation error, specification error, and also ease of use.

2.2.2.1 The Linear Factor Decomposition

In the factor model, asset returns are decomposed linearly into two parts, i.e., the part of returns that is correlated to a set of systematic risk factors and the part of asset-specific returns:

$$\mathbf{r} = \boldsymbol{\alpha} + \mathbf{Bf} + \boldsymbol{\varepsilon} \tag{2.22}$$

where $\mathbf{r}$ is an *n*-dimensional vector of excess returns $\mathbf{r} = (r_1, r_2, ..., r_n)'$, $\mathbf{f}$ is a vector of *k* common risk factors $(k < n)$, $\mathbf{B}$ is an $n \times k$ matrix of factor loadings, and $\boldsymbol{\varepsilon}$ is an *n*-vector of asset-idiosyncratic returns. The vector of coefficients $\boldsymbol{\alpha}$ is set so that $E(\boldsymbol{\varepsilon}) = 0$. The residuals $\boldsymbol{\varepsilon}$ are assumed to be uncorrelated with the factors $\mathbf{f}$. The covariance matrix can thus be represented as:

$$\Sigma = \mathbf{B} \Sigma_{\mathbf{f}} \mathbf{B}' + \Sigma_{\boldsymbol{\varepsilon}} \tag{2.23}$$

where $\Sigma_{\mathbf{f}}$ is the covariance matrix of the factors, and $\Sigma_{\boldsymbol{\varepsilon}}$ is the covariance matrix of the asset-specific returns.

In the *strict factor* model suggested by Sharpe (1963) and Ross (1976), the asset-specific returns are assumed to be cross-sectional uncorrelated $E(\varepsilon_i, \varepsilon_j) = 0$ for $i \neq j$, and so the covariance matrix $\Sigma_{\boldsymbol{\varepsilon}}$ is a diagonal matrix $\Sigma_{\boldsymbol{\varepsilon}} = diag\left\{\sigma_{\varepsilon_i}^2\right\}$. Chamberlain (1983), however, finds the uncorrelated residual assumption unnecessarily strong and suggests an alternative *approximate factor* model, in which idiosyncratic components are allowed to be weakly correlated. The approximate factor structure is now widely applied in dynamic factor models.

2.2.2.2 Single Factor Models

The single index model of Sharpe (1963) is an example of single factor models. Sharpe is also the first to advocate using the factor covariance matrix to solve the mean-variance optimisation problem. The single index model has only one systematic risk factor that influences asset returns, i.e., the exposure to the overall movement of the market. In this model, (2.22) is given by

$$\mathbf{r} = \boldsymbol{\alpha} + \boldsymbol{\beta} r_m + \boldsymbol{\varepsilon} \tag{2.24}$$

where r_m is the return on the market portfolio comprising all assets in the market. The market portfolio is normally proxied by a broad market index such as the S&P500 or the FTSE100 index. Returns are now separated into a component that is correlated with the market (the market component) and an uncorrelated component (the residual component). The exposure of each asset to the market portfolio is measured by its market beta β, estimated based on historical data of return on the market portfolio and return on the asset:

$$\beta_i = \frac{\text{cov}(r_i, r_m)}{\text{var}(r_m)}. \tag{2.25}$$

The covariance matrix implied by the single factor model then becomes:

$$\Sigma = \sigma_m^2 \boldsymbol{\beta}\boldsymbol{\beta}' + \Sigma_{\boldsymbol{\varepsilon}} \tag{2.26}$$

where σ_m^2 is the variance of returns on the market portfolio and $\Sigma_{\boldsymbol{\varepsilon}}$ is a diagonal residual covariance matrix with non-zero elements $\sigma_{\varepsilon,i}^2$.

The model significantly reduces the dimensionality of the covariance matrix. We only have to estimate $2n+1$ parameters: n parameters of market betas β_i, n parameters of asset-specific variance $\sigma_{\varepsilon,i}^2$ and the market variance σ_m^2, as compared to $\frac{1}{2}n(n+1)$ parameters in the fully estimated sample covariance matrix. Since more data is available per parameter, we can expect a reduction in estimation error. However, it is likely that a single factor does not fully capture the total covariation among asset returns, and so single factor models may be severely biased and misspecified. The literature has shown that asset returns may be related to factors other than the market returns (see Ross, 1976, Chen et al., 1986, Fama and French, 1992, Fama and French, 1993). For many years, investment professionals have instead relied on multifactor models in portfolio management and risk analysis.

2.2.2.3 Multifactor Models

With more factors incorporated, multifactor models explain asset returns better than single factor models. Multifactor models also provide a more detailed analysis of risk. As the number of factors is normally chosen to be much fewer than the number of assets, multifactor models reduce the dimensionality of the covariance matrix, making them convenient for financial application. If the factors are uncorrelated, there are only

$k + n(k+2)$ parameters to be estimated, compared to $\frac{1}{2}n(n+1)$ parameters in the sample covariance matrix. For example, for a 25-asset portfolio with three factors, we need to estimate only 78, not 325, parameters. This section presents an overview of the three popular types of multifactor models in practice.

Statistical Factor Models

In statistical factor models, statistical techniques, such as Principal Components Analysis (PCA) or factor analysis, are used to extract the most important uncorrelated sources of variations in asset returns. Statistical factor models use a few linear combinations of returns, or components/factors, which capture most of the variations present in asset returns, to explain the structure of the covariance matrix.

The PCA technique can be applied to either the covariance matrix or the correlation matrix. Here I illustrate the application of the PCA technique to the covariance matrix. Note that since the correlation matrix is just the covariance matrix of the standardised return vector $\mathbf{r}^* = \mathbf{D}^{-1}\mathbf{r}$, where $\mathbf{D}$ is the diagonal matrix of the standard deviations of returns, the application of the PCA to the correlation matrix is straightforward to derive. The PCA analysis employs the eigenvector-eigenvalue decomposition for the symmetric positive semi-definite covariance matrix Σ of returns:

$$\Sigma = \mathbf{V\Lambda V}' \tag{2.27}$$

where $\mathbf{\Lambda}_{n\times n}$ is a diagonal matrix of eigenvalues λ of Σ, ordered from the highest to the lowest, and $\mathbf{V}$ is an $n\times n$ matrix of eigenvectors $\mathbf{v}$ of Σ. $\mathbf{v}_m = [v_{1m}, v_{2m}, ..., v_{nm}]'$ is the eigenvector corresponding to the eigenvalue λ_m. Define P_m as the m^{th}-component of the system. P_m can be represented as a linear combination of returns:

$$P_m = v_{1m}r_1 + v_{2m}r_2 + ... + v_{nm}r_n = \sum_{i=1}^{n} v_{im}r_i \tag{2.28}$$

where r_i is return on asset *i*, or in matrix terms:

$$P_m = \mathbf{v}_m'\mathbf{r}\,. \tag{2.29}$$

Putting together all the components of the system, we get:

$$\mathbf{P} = \mathbf{V}'\mathbf{r}\,. \tag{2.30}$$

The variances and covariances of the components are given by

$$\mathbf{P}\mathbf{P}' = \mathbf{V}'\mathbf{r}\mathbf{r}'\mathbf{V} = (T-1)\mathbf{V}'\Sigma\mathbf{V} = (T-1)\mathbf{V}'\mathbf{V}\mathbf{\Lambda}\mathbf{V}'\mathbf{V} = (T-1)\mathbf{\Lambda}, \tag{2.31}$$

as the eigenvectors are orthogonal ($\mathbf{V}\mathbf{V}' = \mathbf{I}_n$). Since $\mathbf{\Lambda}$ is a diagonal matrix, the components $\mathbf{P}$ are uncorrelated $E(P_i, P_j) = 0$ $(\text{for}\, i \neq j)$ and the variance of the m^{th}-component P_m is the corresponding eigenvalue λ_m. The proportion of the total variations in the covariance matrix that is explained by the m^{th}-component is $\lambda_m / \sum \lambda$. In the correlation matrix, as the sum of the eigenvalues is equal to the number of the eigenvalues, the explanatory power of the m^{th}-component on the correlation is measured by λ_m / n. Since the eigenvalues are ordered according to size, the first principal component will explain the greatest amount of the total variations in the covariance matrix.

Because the eigenvectors are orthogonal $\mathbf{V}\mathbf{V}' = \mathbf{I}_n$, (2.30) can be rewritten as $\mathbf{r} = \mathbf{V}\mathbf{P}$. Hence each asset return can be represented as a linear combination of the components:

$$r_i = v_{i1}P_1 + v_{2m}P_2 + \ldots + v_{ni}P_n. \tag{2.32}$$

Choose the first k components $\mathbf{P}^* = (P_1, P_2, \ldots, P_k)$ that explain the most part of the variations of asset returns, then $\mathbf{r} = \mathbf{V}^*\mathbf{P}^* + \boldsymbol{\varepsilon}$, where $\boldsymbol{\varepsilon}$ is the vector of residuals, accounting for the remaining small variations that are not explained by the first k components $\mathbf{P}^*$. As the components are uncorrelated, the covariance matrix will then be reproduced with fewer factors:

$$\Sigma = \mathbf{V}^*\mathbf{\Lambda}^*\mathbf{V}^{*'} + \Sigma_\varepsilon \tag{2.33}$$

where $\mathbf{\Lambda}^*$ is the diagonal matrix of the components' variances $\mathbf{\Lambda}^* = diag(\lambda_1, \lambda_2, \ldots, \lambda_k)$, $\mathbf{V}^*$ is the component loadings and Σ_ε is the diagonal covariance matrix of the residuals

The PCA technique is very useful in highly correlated systems, where most of covariations can be explained by just a few independent sources. The disadvantage of this technique nevertheless lies in the interpretation. The factors extracted from the PCA are statistical artifacts, and so practitioners have to interpret the economic meaning of these statistically-derived factors. The interpretability of the factors is crucial in determining the validity of the PCA.

Macroeconomic Factor Models

Macroeconomic factor models trace the sources of the common covariations among assets to observable macroeconomic variables, such as inflation, interest rates and business cycle. Unlike statistical factor models, this approach is backed economically since asset returns are systematically affected by macroeconomic conditions. Practitioners observe and study the magnitude and persistence of macroeconomic variables in explaining historical stock returns and choose pervasive factors for the models. A good factor must be able to explain the covariations of asset returns, as well as be easy to interpret and robust over time. The factor betas of each asset are typically estimated in time series regressions of asset returns on the given factors. Chen et al. (1986) are among the first to explore and test macroeconomic factors that affect the behaviour of stock returns. They establish a set of five economic variables that can affect the discount rate and/or future dividends, which in turn influence the stock prices in the US market. These include growth in industrial production, changes in expected inflation, unexpected inflation, unexpected changes in risk premia, and unexpected changes in the term structure slope. The five-factor Chen, Roll and Ross (CRR) model has been popularly applied and extended in many researches. A notable example is the BIRR (Burmeister, Ibbotson, Roll, and Ross) model that employs a similar list of five macroeconomic factors in the APT framework to construct superior portfolios (see Burmeister et al., 2003). The CRR model is nevertheless subject to criticism. For example, Shanken and Weinsten (2006) challenge the robustness of the CRR factors. They design an experiment that is comparable to that of the CRR model in most respects, except only for the use of pre-ranking returns to estimate betas. Using instead the post-ranking returns as in Fama and French's study (1992), they claim that there is no indication that the macroeconomic factors, except industrial production, are priced.

Fundamental Factor Models

Fundamental factor models concentrate on the explanatory power of security attributes, such as market capitalization, industry, book-to-price ratio, dividend yield, on stock returns. These security characteristics have been found to be surprisingly powerful in describing the comovement of individual equities. The fundamental factor betas, unlike those of macroeconomic factors, may not need to be estimated from time series regressions. Fundamental factor models may use a company's observed attributes such as firm size, industry classification as betas. In this sense, factor betas are exogenously determined, firm-specific attributes rather than estimated sensitivities to random factors (Connor, 1995). For instance, the industry factor beta is a dummy variable, taking the value of one if a firm belongs to an industry and of zero if it does not.

One of the most popular fundamental factor models is the Fama-French three-factor model. Fama and French (1992, 1993) suggest that two security attributes, market capitalization and book-to-price ratios, are strongly correlated with the difference in mean returns across securities; smaller capitalization and higher book-to-price ratio stocks are found to have higher mean returns. Arguing that these higher mean returns are due to non-diversifiable portfolio risk, Fama and French form two combination portfolios, i.e., the SMB (small minus big) and the HML (high minus low) portfolios and use these two portfolios, together with the value-weighted market index as the pervasive risk factors. Factor betas are estimated by time series regressions of the asset returns on returns of the factor portfolios. Carhart (1997) and Jegadeesh and Titman (2001) extend Fama and French's three-factor model, adding an additional factor-mimicking portfolio to represent the momentum factor (proxied by high-twelve-month returns minus low twelve-month-returns). Goyal and Santa-Clara (2003) and Ang et al. (2006, 2009) also suggest evidence of an own-volatility-related factor that adds explanatory power to the Fama-French model, for explaining both return comovements and mean returns. Commercially, MSCI BARRA has developed a multifactor model covering the world's major equity markets. For instance, for the US market, their model consists of 12 risk indices such as volatility, size, growth, earning-to-price, book-to-price, financial leverage and 55 industry dummies, further classified into 13 industry categories.

2.2.2.4 Practical Implementation and Issues

Factor models have gained significant popularity in practical portfolio management. The parsimony of the factor structure reduces the number of estimated parameters, and hence reduces estimation error. Besides, the factor structure also avoids the ill-conditioned problem of the inverse covariance matrix, providing a better conditioned alternative to the fully estimated covariance matrix (Fan et al., 2008). However, to reduce the dimensionality of the covariance matrix, the number of factors is normally chosen to be much smaller than the number of assets $(k << n)$, which means factor models could concentrate only on the strongest sources of covariations. Thus, the models may be mispecified in the sense that they omit some important sources, or that the chosen factors may be transitory and lose their explanatory power over the next period. Increasing the number of factors offers more flexibility in approximating the data generating process, but at a cost of estimation error. Besides, including too many factors may run the risk of overfitting the data, producing poor out-of-sample forecast performance.

Another vital concern in applying the multifactor models is to decide which types of factors to include. Connor (1995) compares the explanatory power of the three types of multifactor models for US equity returns. Theoretically, he shows that the three types are not necessarily inconsistent; in the absence of estimation error and with no limits on data availability, the three models are simply restatements or rotations of one another. However, in practice where estimation error is common and data is restricted, the three models may differ. Connor finds empirically that the fundamental and statistical multifactor models outperform the macroeconomic factor models in terms of explanatory power. He also finds that the explanatory power of a macroeconomic factor model is marginally negligible when it is added to a fundamental factor model, implying that fundamental factors may capture the same risk as the macroeconomic factors. On the contrary, Burmeister et al. (2003) advocate the use of macroeconomic factors while raising concerns about the fundamental factor models. They argue that the fundamental factors are based on accounting data that may come from different accounting rules, or even if they are from the same accounting rules they may be released at different report dates that makes it difficult to obtain time-synchronised comparison. Nevertheless, macroeconomic factor models incur the same problem. Economic variables such as GDP, inflation are normally released at different time, and since they are aggregate variables, their estimation error may also be very large. Statistical factor models, though estimated by maximum explanatory power, are faced with different problems. After extracting factors from the PCA or factor analysis, practitioners have to interpret the economic meaning underlying those factors. To make thing more complicated, factors change over time, which means a third factor in one sample period may be completely different from the third factor in another sample period.

Determining the number of factors is also a central issue, especially in statistical models. The factors, of course, should be robust, statistically significant, and justified by an economic intuition. Until lately, the number of factors in statistical models was often assumed rather than determined by the data. Fortunately, recent studies have proposed tests to estimate the number of factors from observed data. Connor and Korajczyk (1993) suggest a test under a sequential limit asymptotic assumption, i.e., n converges to infinity with a fixed T and then T converges to infinity. Starting with some certain factors, they will add an additional factor if they find an increase in explanatory power from adding that factor. Their tests are based on the difference in the cross-sectional averages of asset-specific variances with and without the additional factor. Bai and Ng (2002) have a different approach. They set up the determination of factors as a model selection problem. Working in the PCA framework, they develop two types of information criterion, which represent the trade-off between good fit and parsimony. Their criteria are developed under the assumption that $n, T \to \infty$, thus applicable for many datasets. Another approach is motivated by the work on Random Matrix Theory for stock correlations of, e.g., Plerou et al. (2002). Benghtsson and Holst (2002) suggest choosing the number of principal components that are determined by the eigenvectors corresponding to the eigenvalues that deviate significantly from the maximum eigenvalue bound obtained for a random matrix.

In portfolio choice and risk management, factor models have been popularly applied to produce better estimates of the covariance matrix. Chan et al. (1999) study the performance of different factor models in a portfolio choice problem. Testing the predictive power of different factor models (ranging from one factor to ten factors), they show that factor models clearly improve the forecast performance of the covariance matrix. However, they also find that only a few factors such as the market, size, book-to-market value of equity are sufficient in capturing the general structure of the covariance matrix. Extending the number of factors beyond this relatively small set does not lead to superior forecast performance. In another study, Briner and Connor (2008) compare performance of three covariance matrix estimators, i.e., the sample covariance matrix, the single market factor model, and the multifactor model. Their simulation and empirical results show that the multifactor model performs best for large investment universes and typical sample lengths. This result is consistent with conventional wisdom, proposing that the market model underperforms because of excessive specification error, while the sample covariance matrix underperforms due to high estimation error.

2.2.3 Shrinkage Models

Ledoit and Wolf (2003, 2004) extend the mean shrinkage estimator of James and Stein (1961) to the covariance matrix. They propose a shrinkage covariance matrix estimator $\hat{\Sigma}_s$ that is a convex combination of the usual sample covariance matrix $\hat{\Sigma}$ and a shrinkage target $\mathbf{S}$ (or its estimate $\hat{\mathbf{S}}$):

$$\hat{\Sigma}_s = \alpha \hat{\mathbf{S}} + (1-\alpha)\hat{\Sigma} \tag{2.34}$$

where α is the shrinkage constant or shrinkage intensity, $\alpha \in [0,1]$. The underlying idea is to shrink the sample covariance matrix to the shrinkage target so as to address the trade-off between estimation error and specification error. The sample covariance matrix is unbiased but full of excessive estimation error, while the shrinkage target, due to their simple dimensionality, has less estimation error, but may be misspecified. The new covariance matrix $\hat{\Sigma}_s$ can be seen as a weighted average of the biased and unbiased estimators with the weight α. This weight α, representing the optimal trade-off, controls how much structure to be imposed: the heavier the weight, the stronger the imposed structure.

As suggested by Ledoit and Wolf, the shrinkage target should fulfil two requirements: (i) involving a small number of parameters, and (ii) reflecting important characteristics of the true covariance matrix. Ledoit and Wolf (2003, 2004) choose the single-index factor model of Sharpe (1963) and the constant correlation model of Elton and Gruber (1973) as their shrinkage targets. Bengtsson and Holst (2002) extend the study of Ledoit and Wolf to shrink the covariance matrix to a *k*-factor model derived from a PCA analysis. A positive definite target also guarantees the positivity of the shrinkage estimate, even when the sample covariance matrix itself is singular. This makes shrinkage a particularly practical statistical tool for constructing large-scale equity portfolios.

Ledoit and Wolf develop algorithms to estimate the shrinkage constant α by minimising a loss function that does not involve the inverse of the covariance matrix. This is an advance as previous shrinkage intensity estimation was normally based on loss functions involving the inverse covariance matrix, which makes the estimators break down when $n > T$. In their two articles, Ledoit and Wolf apply their shrinkage covariance matrix estimator in a global minimum variance portfolio choice problem and show that their estimator produces portfolios with significantly lower out-of-sample variance than those produced by a set of well-established competing approaches, including the multifactor models. Interestingly enough, they find that the shrinkage intensity with a single-factor target is remarkably stable through time with a value of around 0.8, suggesting that there is about four times as much estimation error in the sample covariance matrix as there is bias in the single-factor covariance matrix. Janagathan and Ma (2000), however, challenge the complicated algorithm of Ledoit and Wolf to estimate α. They argue that different covariance matrix estimators contain error in different directions, hence using a portfolio of covariance matrix estimators (e.g., a simple average of a sample covariance matrix, a single index estimate, and a matrix consisting of the diagonal part of the sample covariance matrix) may cancel the error out. Motivated by the study of Jagannathan and Ma, Disatnik and Benninga (2007), while acknowledging that the shrinkage estimator of the covariance matrix is indisputably better than the sample covariance matrix estimator, find no statistical differences in the ex post standard deviations of the global minimum variance portfolios constructed with the more sophisticated shrinkage estimator of Ledoit and Wolf and those with simpler portfolios of estimators of Jagannathan and Ma.

2.2.4 The Constant Correlation Coefficient Model

Given the drawbacks of the sample covariance matrix as an input to the portfolio optimiser, Elton and Gruber (1973) suggest the use of a constant correlation coefficient model, where all pairwise correlation coefficients are assumed to be equal and equal to the average pairwise correlation coefficient. In the empirical study, they show that their estimator is both statistically significant in producing better five-year estimates of future correlation coefficients, and economically significant in yielding superiour out-of-sample portfolio performance, even with or without short sales constraints, than those produced from the sample or single factor covariance matrices.

Arguing that the model of Elton and Gruber still has many parameters to estimate (all the pairwise correlation coefficients have to be estimated to obtain their average), Aneja et al. (1989) suggest a simplified but exact portfolio approach of estimating the average correlation coefficient without having to estimating all the pairwise correlations. To estimate the covariance matrix of an *n*-asset portfolio, their approach will only have to estimate $n+1$ variances: the variances of n assets and the variance of a portfolio where investment in each security equals to the reciprocal of its sample standard deviation.

This method greatly reduces the computational requirements for estimating the average correlation coefficient.

2.3 Conclusion

Markowitz's mean-variance optimisation theory provides a convenient and objective framework to allocate wealth in a portfolio. The theory also beautifully captures the two fundamental economic insights of risk-return trade-off and diversification. However, despite its obvious appeal, the Markowitz paradigm is faced with several criticisms. Theoretically, academics have consistently attacked on the prohibitively restrictive assumptions of the Markowitz analysis, e.g., the quadratic utility, the single-period investment horizon. In application, practitioners have traditionally resisted the use of the classical framework, not least because of the difficulty in estimation of the inputs. The limitations of Markowitz theory have spurred numerous extensions of the paradigm. Many models have been suggested, ranging from small calibrations of estimation of the moments of returns to incorporating sophisticated statistical developments into the optimiser. Each approach has its own advantages and limitations that make no approach emerge as a clear favourite. However, in a striking study, DeMiguel et al. (2009) show that the naïve diversified equally weighted portfolio cannot be dominated by any of the fourteen popularly used optimal portfolio models, ranging from the classical sample mean-variance efficient strategy to the Bayesian approach to estimation error, or to the models that impose constraints on portfolio weights. Despite considerable progress in the design of optimal portfolios, estimation error in expected returns and the covariance matrix may still be so excessive that it erodes all the gains from optimal, relative to naïve diversification. The needs for more reliable estimates of the moments of asset returns still pose significant challenges. Exploiting the predictability of the covariance matrix in conditional volatility models suggests an interesting direction. The next chapter will turn to the analysis of conditional covariance matrix estimators and their implications for practical asset allocation.

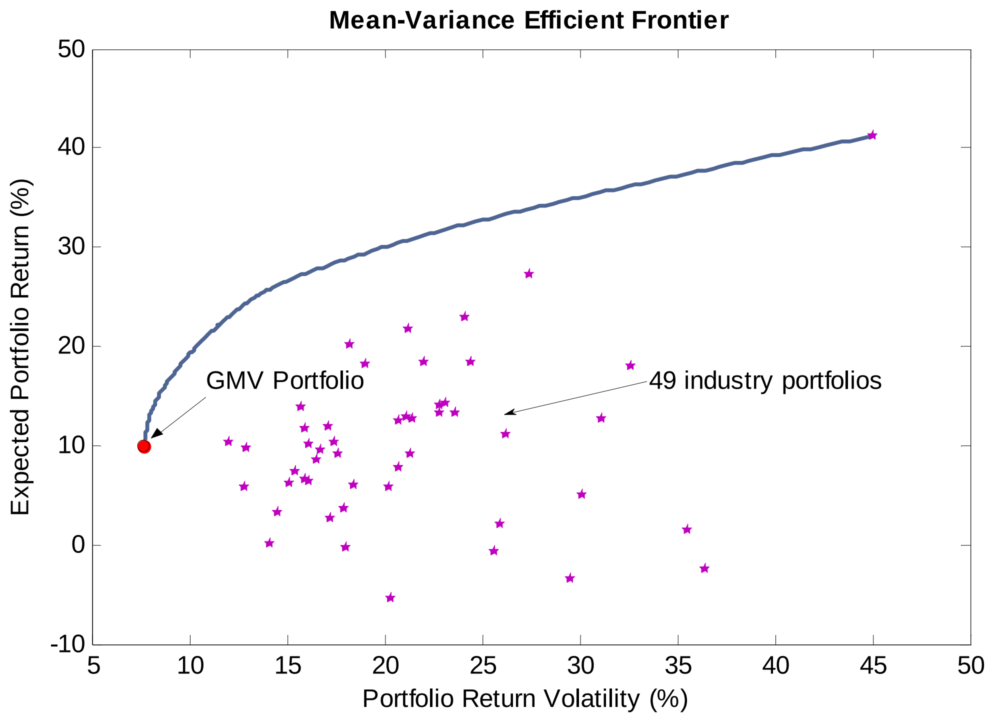

Figure 2.1. The Mean-Variance Efficient Frontier.

The figure plots the mean-variance efficient frontier of 49 average-value-weighted industry portfolios of the US, using data of monthly returns. Expected returns and volatilities are annualised.

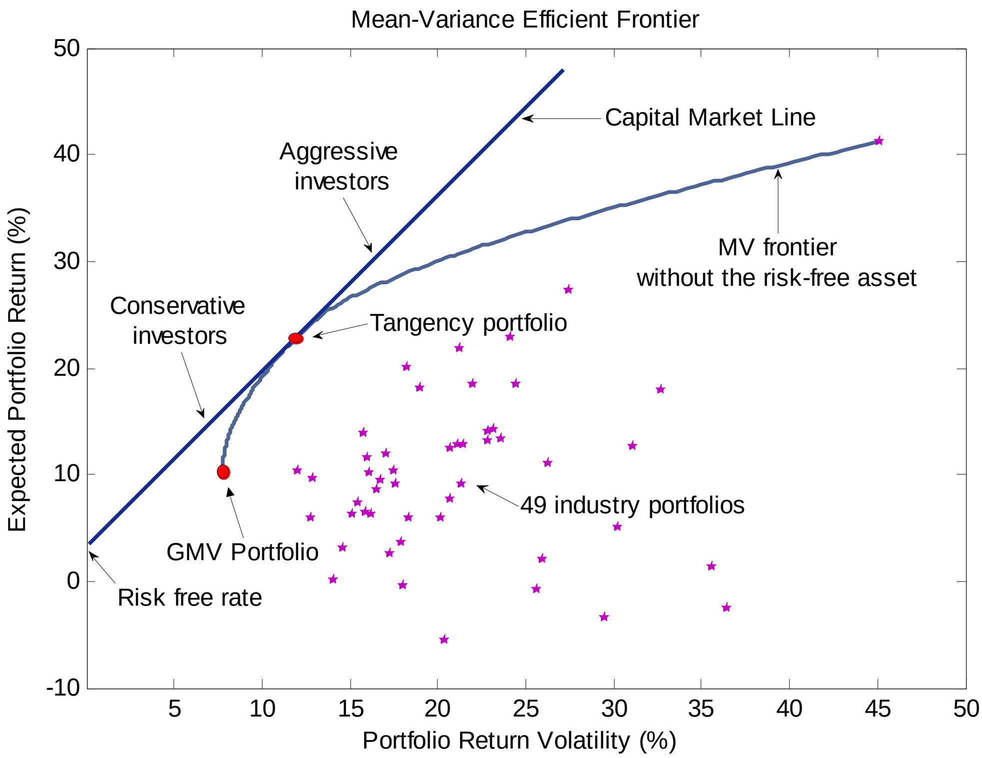

Figure 2.2. The Mean-Variance Efficient Frontier and The Capital Market Line.

The figure plots the mean-variance efficient frontier and the Capital Market Line, constructed from 49 average-value-weighted industry portfolios of the US, using data of monthly returns. The risk-free rate is assumed 4%. Expected returns and volatilities are annualised.

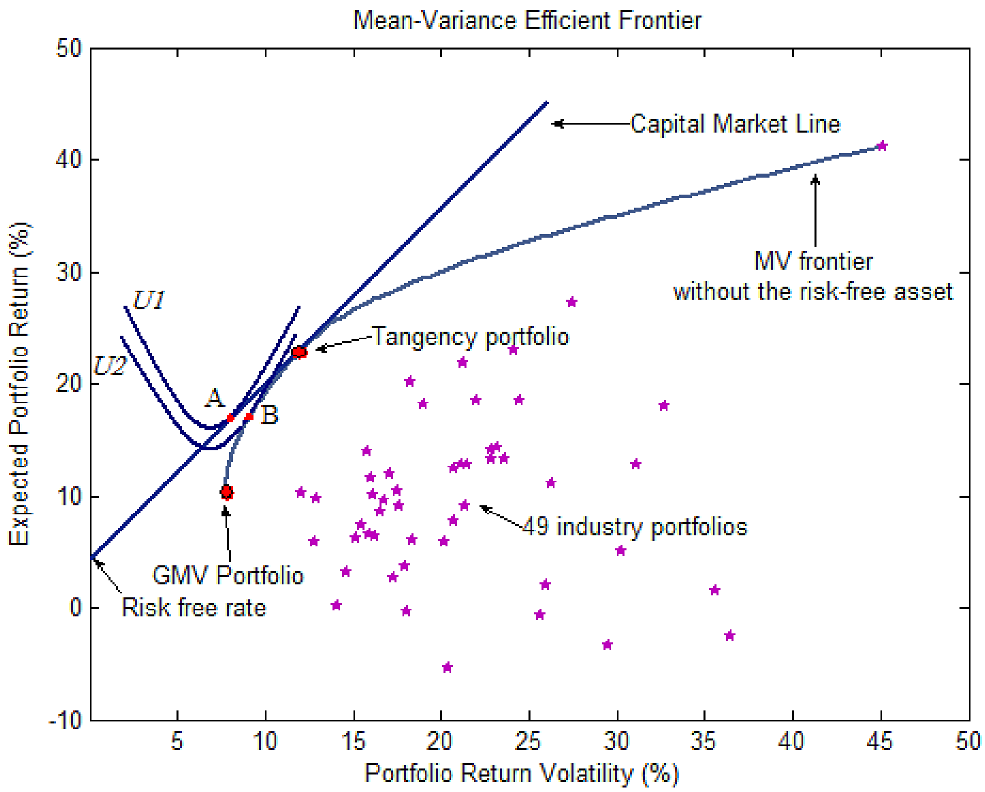

Figure 2.3. Utilities and Optimal Portfolios.

The figure shows how an investor chooses his optimal portfolios based on the efficient frontiers and his utility indifference curves. The risk-free rate is assumed 4%. Expected returns and volatilities are annualised.

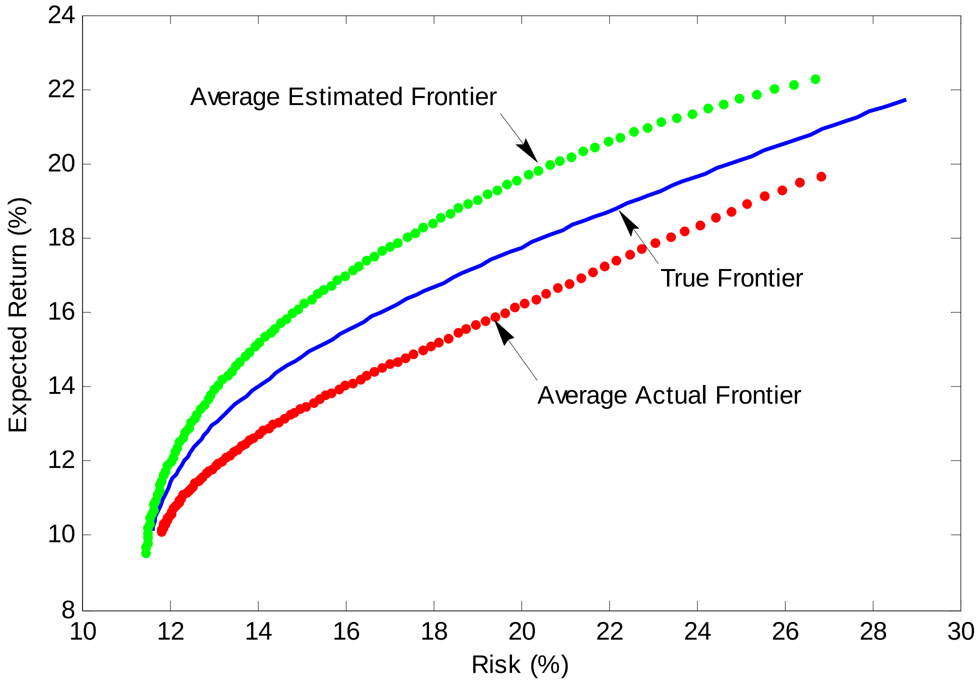

Figure 2.4. The True, Estimated and Actual Mean-Variance Efficient Frontiers.

The figure plots the mean-variance frontiers generated using real and simulated data of 10 US industry portfolios (T = 120 observations for the simulated data). The true and the estimated frontiers are constructed using the true and the estimated parameters, respectively. The actual frontiers are obtained by applying portfolio weights derived from the estimated parameters to the true parameters to calculate portfolios' expected returns and risk. Expected returns and volatilities are annualised.

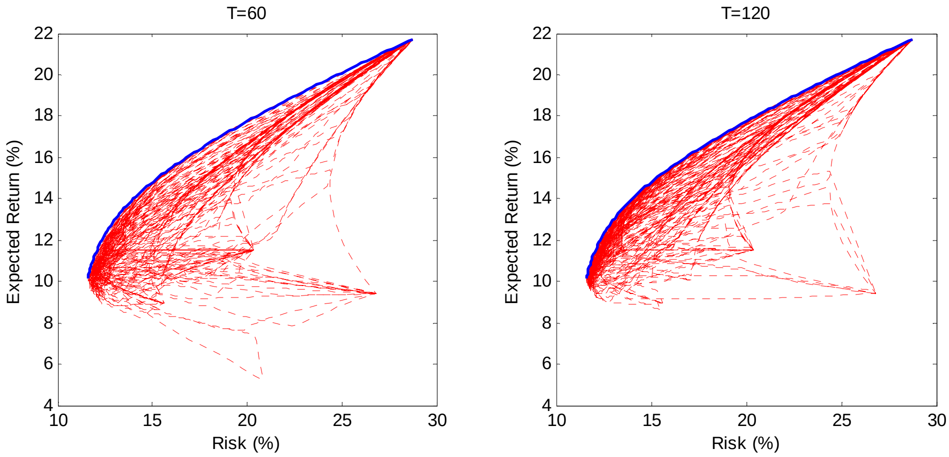

Figure 2.5. The True and Actual Mean-Variance Efficient Frontiers.

The figure illustrates the estimation error problem in using the sample estimates of expected returns and the covariance matrix to construct the mean-variance efficient frontiers. The blue solid curve is the true efficient frontier, while the red dashed curves are the 250 simulated actual frontiers. The frontiers are constructed using real and simulated data of 10 US industry portfolios for two sample sizes of 60 and 120 monthly observations. Expected returns and volatilities are annualised.

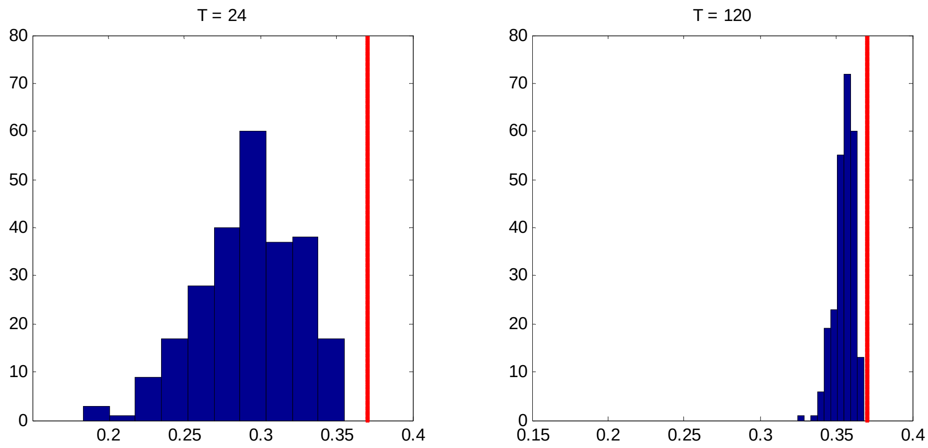

Figure 2.6. The Sharpe Ratios of the Tangency Portfolios.

The figure shows the histograms of the Sharpe ratios of the 250 actual tangency portfolios. The portfolios are constructed using real and simulated data of 10 US industry portfolios with two sample sizes of 24 and 120 monthly observations. The investor is assumed to know the true expected returns but uses the sample estimates of the covariance matrix. The vertical red line in each plot represents the Sharpe ratio of the true tangency portfolio.

Chapter 3

The Time-Varying Conditional Variance-Covariance Matrix

Applications of the classical mean-variance portfolio optimisation typically assume a constant return distribution, in which expected returns and risk do not change over time. However, it is now well established that the covariance matrix of short horizon financial asset returns is both time-varying and highly persistent. Starting from the seminal work of Engle (1982), a number of conditional volatility models such as the multivariate Exponentially Weighted Moving Average (EWMA), multivariate Generalised Autoregressive Conditional Heteroskedasticity (GARCH) and multivariate Stochastic Volatility (SV) models have been developed to capture these features of the covariance matrix. These models are now routinely used in many areas of applied finance, including asset allocation, risk management and asset pricing. Mounting evidence now suggests that multivariate conditional volatility models produce better forecasts of the covariance matrix than those produced by the unconditional covariance matrix estimator (see, for example, Engle and Colacito, 2006). Practical problems, such as asset allocation, consequently benefit from better forecasts of the covariance matrix. Indeed, ample evidence clearly demonstrates that dynamic asset allocation strategies, based on time-varying conditional covariance matrices, dominate static strategies, based on constant unconditional alternatives (see, for example, Fleming et al., 2001, Han, 2006, Thorp and Milunovich, 2007).

Estimation of the time-varying conditional covariance matrix has been the subject of extensive research. This chapter provides a summary of some popular conditional volatility models and their application to asset allocation. The chapter will primarily focus on the Moving Average and the GARCH models, highlighting their similarities and differences. Stochastic Volatility, Realised Volatility, Option Implied Volatility models are out of the research scope of this chapter and of the thesis. Due to space limits, the chapter will not provide an exhaustive list of all Moving Averages and GARCH models, nor cover all details of each model. Particularly, I will not cover some areas, such as the testing and estimation procedure, or the forecast evaluation. The purpose here is to help readers get a glimpse of the developments of conditional volatility models in both univariate and multivariate context, before guiding them through the multivariate long memory conditional volatility models studied in the next chapters. More details of the conditional volatility models can be found in excellent reviews of, among others, Poon (2005), Andersen et al. (2006), Xiao and Aydemir (2007).

Section 3.1 begins with some well-known properties of asset return volatility. Extensive research that captures these properties to produces better estimates and forecasts of the covariance matrix is then summarised in the following sections. My focus is especially on the multivariate conditional volatility models that are applicable to a large number of assets. Section 3.2 presents the Moving Average models, with due attention paid to the widely-used Riskmetrics Exponentially Weighted Moving Average (EWMA) of JP Morgan (1994). Section 3.3 is devoted to the GARCH family. Alternative univariate GARCH models are discussed in Section 3.3.1, while their multivariate generalisations are detailed in Section 3.3.2. Owing to their importance in the research, long memory conditional volatility models are studied in a separate section (Section 3.4). Finally, Section 3.5 briefs some applications of the conditional volatility models in the asset allocation framework.

3.1 Properties of Asset Return Volatility

This section introduces some of the well-established properties of asset return volatility. The recognition of those properties has sparked off the development of numerous conditional volatility models in the last 30 years.

Fat tails

The unconditional distribution of asset returns is known to exhibit fatter tails than those exhibited in the normal distribution. An ample body of evidence suggests that although the normal distribution may closely explain financial asset returns in the middle of the curve where most gains and losses occur, it fails to do so in the extreme edges. There are more days of spectacular price increases or falls than it is expected under the normal distribution assumption. For example, we typically observe financial returns of four standard deviations many days in a year, which is inconsistent with the normal distribution. This non-normality feature of the asset return distribution should be taken into account in the construction of any volatility model.

Volatility persistence

It is generally agreed that volatility is time-varying and persistent. While returns themselves contain little autocorrelation, which is consistent with the efficient market hypothesis, absolute returns and squared returns (proxies of volatility) are found to be highly correlated and persistent. This property holds for returns of equities, bonds, exchange rates, interest rates in different markets and different countries at daily or even weekly frequencies, and is even more pronounced at high frequency intra-day returns.

Volatility persistence is among the first features of volatility to be recognised. Mandelbrot (1963) observes that "large changes tend to be followed by large changes - of either sign- and small changes tend to be followed by small changes." However, Mandelbrot then emphasises the unconditional non-normality of returns, rather than volatility clustering. The first formal study of volatility persistence is credited to Engle (1982), who exploits this feature to develop the Autoregressive Conditional Heteroskedasticity (ARCH) model. Volatility persistence implies that information in the past can be exploited to generate future forecasts of volatility. The seminal work of Engle has served as the foundation for extensive and ongoing research on time-varying conditional volatility.

Mean reversion

Volatility persistence implies that when volatility is high, it is likely to remain high, and vice versa. However, this effect is time bounded so that a period of high volatility will eventually give way to a period of normal volatility and conversely, volatility will rise after a period of low volatility. This 'mean reversion' feature implies volatility will eventually revert to a long-run normal level. Consequently, long-run forecasts of volatility will converge to this normal level, no matter when they are made.

Asymmetric volatility

Volatility asymmetry has been noticed in equity markets. It has been observed that volatility is higher in bear markets than it is in bull markets. A negative return shock (unexpected price drop) will lead to a higher subsequent volatility than a positive return shock (unexpected price increase) of the same magnitude. Black (1976), among others, attributes this phenomenon to the 'leverage effect', in which a fall in stock price increases financial leverage and hence financial risk of the firm, leading to changes in volatility. Pindyck (1984) and French et al. (1987), however, have a different explanation. They argue that the asymmetric nature of volatility response to shocks could simply reflect the time-varying risk premium – the 'volatility feedback'. An anticipated increase in volatility raises the required returns, hence provoking an immediate decline in stock price. Though addressing the same behaviour of volatility, the two approaches have different causality. While the leverage effect treats the return shock as the cause to the conditional volatility, the volatility feedback mechanism treats it as the effect. Which direction dominates has not got a clear-cut answer and still remains an open question for academic researchers.

Long memory behaviour of volatility

A mounting body of empirical evidence now suggests that the autocorrelation function of squared return innovations declines more slowly than the exponential decay implied in the EWMA and GARCH models, and hence volatility shocks are more persistent than these models imply. Ding et al. (1993) are the first to identify the so-called long memory behaviour in volatility. They investigate the volatility of the daily S&P500 index returns and find that the sample autocorrelation function of volatility decreases slowly and remains significantly positive after very long lags; yet, the volatility process is still essentially stationary. This feature is important not only for the measurement of current volatility, but also for forecasts of future volatility, especially over longer horizons.

These properties of financial asset return volatility have spurred the development of numerous volatility models to provide accurate estimates and reliable forecasts of future volatility. In the following sections, various approaches to model the conditional covariance matrix will be investigated.

3.2 Moving Average Models

Moving average models are a simple, yet practically powerful approach to estimate and forecast the time-varying covariance matrix. The simplest specification of this class is the Equally Weighted Moving Average model, in which elements of the covariance matrix are estimated as sample squares and cross products of returns over rolling windows. This is sometimes referred to as 'the historical volatility'. It offers the simplest way of incorporating actual data in the estimation of the time-varying covariance matrix. Another specification is the widely used RiskMetrics Exponentially

Weighted Moving Average (EWMA) model of JP Morgan (1994). This section provides a brief overview of these two models.

3.2.1 The Equally Weighted Moving Average Model

Consider an n-dimensional vector of returns $\mathbf{r}_t = \left(r_{1t}, r_{2t}, \ldots, r_{nt}\right)'$ with conditional mean zero and conditional covariance matrix $\mathbf{H}_t$:

$$\mathbf{r}_t = \mathbf{H}_t^{\frac{1}{2}} \mathbf{z}_t , \tag{3.1}$$

where $\mathbf{z}_t$ is i.i.d with $E\left(\mathbf{z}_t\right) = 0$ and $\operatorname{var}\left(\mathbf{z}_t\right) = \mathbf{I}_n$. The 'historical' covariance matrix is calculated on a T-day window that is rolled through time, each day adding the new return and taking off the oldest return:

$$\mathbf{H}_t = \frac{1}{T} \sum_{i=1}^{T} \mathbf{r}_{t-i} \mathbf{r}_{t-i}^{'} . \tag{3.2}$$

The sophistication of this model lies in the choice of the window length T. If the length is short, the estimate may be noisy since the sampling error is proportional to $1/\sqrt{T}$. The longer the window, the less noisy the estimate, but the more biased it is when far more distant observations, which may not be relevant today, are included in the calculation. Hence, the length of the window T directly determines the trade-off between the sampling error and the unbiasedness of the estimate.

The model captures the time-varying property of volatility and covariance in a simplistic way, through a rolling window. However, by putting equal weights on both recent and distant observations, the model fails to capture the persistence of volatility and covariance. Empirical studies, consequently, suggest that the historical method is not very effective for short-term horizons. Long-term volatility could be estimated with this method, but only when we assume that the past is an accurate reflection of the future.

3.2.2 The Exponentially Weighted Moving Average Model

Unlike the Equally Weighted Moving Average model, the Exponentially Weighted Moving Average (EWMA) model puts more weight on the recent observations and less on the distant past, hence capturing the volatility persistence and enabling volatility to react faster to shocks. The impact of the shocks also dies out exponentially instead of remaining the same until they are excluded out of the equally weighted model. The EWMA covariance matrix has the following specification:

$$\mathbf{H}_t = (1-\lambda)\mathbf{r}_{t-1}\mathbf{r}_{t-1}^{'} + \lambda\mathbf{H}_{t-1}, \tag{3.3}$$

where λ is the decay factor $(0 < \lambda < 1)$. The first term of the right hand side of (3.3), $(1-\lambda)\boldsymbol{r}_{t-1}\boldsymbol{r}_{t-1}'$, denotes the response of volatility to one-period news, while the second term, $\lambda\mathbf{H}_{t-1}$, determines the persistence in volatility. The higher the value of λ, the more persistent the process and the slower the response to new shocks. However, in the EWMA model, the reaction and persistence parameters are not independent because they sum to one.

By backward substitution, the covariance matrix can be written as:

$$\mathbf{H}_t = (1-\lambda)\sum_{i=1}^{\infty}\lambda^{i-1}\mathbf{r}_{t-i}\mathbf{r}_{t-i}^{'}. \tag{3.4}$$

The model derives its name from this formulation, in which the elements of the covariance matrix are the exponentially weighted moving averages of past squares and cross products of returns. In practice, the process is often estimated with a cut-off time *T*, scaling the infinite sum in (3.4) by

$$\mathbf{H}_t = \frac{(1-\lambda)}{1-\lambda^T}\sum_{i=1}^{T}\lambda^{i}\mathbf{r}_{t-i}\mathbf{r}_{t-i}^{'}. \tag{3.5}$$

Under the RiskMetrics (1994), λ takes the values of 0.94 and 0.97 for daily and weekly forecasts, respectively. The EWMA process is hence easily estimated in a spreadsheet for any dimensional system. The one-step ahead forecast is readily given in the model:

$$\mathbf{H}_{t+1} = (1-\lambda)\mathbf{r}_t\mathbf{r}_t^{'} + \lambda\mathbf{H}_t. \tag{3.6}$$

By recursive substitution, the *h*-step forecast is equal to the one-step ahead forecast:

$$\mathbf{H}_{t+h} = \mathbf{H}_{t+h-1} = \ldots = \mathbf{H}_{t+2} = \mathbf{H}_{t+1}. \tag{3.7}$$

Assuming returns are serially uncorrelated, the expected covariance matrix over *k* cumulative steps is given by $\mathbf{H}_{t+1:t+k} = k \times \mathbf{H}_{t+1}$. The multiple-period forecast is a simple product of the one-day forecast with the forecast horizon, *k*. This is also known as the 'square root of time' rule for volatility forecasts. The EWMA model can thus be thought of as a random walk model, where a shock will have a permanent effect on the expectation of future variance and covariance. The volatility process in the EWMA model is not mean-reverting, which is quite counterfactual since financial return volatility tends to eventually converge to its long-run average.

The multivariate Riskmetrics EWMA model of JP Morgan (1994), though being non-mean-reverting and very restrictive when imposing the same degree of smoothness on all elements of the estimated covariance matrix, enjoys the most popular practical application among multivariate conditional volatility models due to its high parsimony.

3.3 GARCH Models

Observing that squared residuals are often autocorrelated even though residuals themselves are not, Engle (1982) sets the stage for the new class of time-varying conditional volatility models with the Autoregressive Conditional Heteroskedasticity (ARCH) model. The new model has inspired a huge amount of related research on its development, generalisation and application, and deserved Engle a Nobel Prize in Economics in 2003.[3] This section introduces the ARCH model and some of its popular generalisations in both univariate and multivariate context.

3.3.1 Univariate GARCH Models

3.3.1.1 The Basic ARCH Model

The ARCH model of Engle (1982) parallels the Wold representation for the conditional mean to modelling the conditional variance. Engle is the first to treat the unconditional mean and variance as constant, while letting both the conditional mean and variance be time-varying. Allowing for the time-varying conditional variance (conditional heteroskedasticity), the ARCH model successfully captures the persistent volatility feature of financial time series, providing a natural and powerfully simple framework for estimating and forecasting volatility.

[3] Engle shared the 2003 Nobel Prize in Economics with Granger. Engle's contribution was recognised "for methods of analyzing economic time series with time-varying volatility (ARCH)," whereas Granger's was "for methods of analyzing economic time series with common trends (cointegration)."

Let r_t be the log return of an asset at time *t*. The return is decomposed into an expected conditional mean $\mu_t = E\left(r_t \middle| \mathcal{F}_{t-1}\right)$ based on the information set $\mathcal{F}_{t-1}$ available at time $t-1$ and an innovation ε_t. The return series has no serial autocorrelation or minor autocorrelation, if any.[4]

$$r_t = \mu_t + \varepsilon_t \tag{3.8}$$

$$\varepsilon_t = \sqrt{h_t} z_t, \tag{3.9}$$

where z_t is a white noise process with zero mean and unit variance and h_t is the conditional variance at time *t*. In practice, z_t is often assumed to follow the standard Gaussian or the standardised Student-*t* distributions. In the ARCH model, the residuals ε_t are serially uncorrelated while their squares are autocorrelated over time. In the following, to facilitate the presentation, the conditional mean is assumed constant and equal to zero, a common assumption in risk management at least when a short horizon is considered. Under the ARCH*(p)* model, the conditional variance is estimated by taking the weighted average of past squared errors:

$$h_t = \omega + \sum_{i=1}^{p} \alpha_i \varepsilon_{t-i}^2 \tag{3.10}$$

with $\omega > 0$ and $\alpha_i \geq 0$ to ensure the strict positivity of σ_t^2. Under this structure, large past changes (large ε_{t-i}^2) imply that the current conditional variance h_t is also large, and vice versa. The ARCH model is thus able to capture the volatility clustering observed in asset returns. One advantage of the ARCH model is that the weight α_i can be estimated from historical data, based on, e.g., the Maximum Likelihood procedure, even though the 'true' volatility is never observed. The unconditional variance of r_t is $\sigma^2 = \omega \big/ \left(1 - \sum \alpha_i\right)$, a constant even though the conditional variance is time-varying. The ARCH*(p)* model is covariance stationary if $\sum \alpha_i < 1$.

[4] The conditional mean equation can be of any form. However, as the ARCH family concentrates on modelling the conditional variance, they usually have a simple conditional mean equation to extract all serial autocorrelations in the residuals. Many of the ARCH models in practice just let the simple conditional mean to follow a stationary ARMA process or even assume that μ_t is equal to zero.

A good volatility model must produce good forecasts. As σ_t^2 is serially correlated, this dependence can be exploited to produce accurate volatility forecasts. From (3.10), h_t is known at time $(t-1)$, so the one-step forecast is readily available. The multi-step forecast can be formulated by assuming that $E\left(\varepsilon_{t+h}^2\right)=h_{t+h}$.

Another remarkable feature of the ARCH model is that the implied unconditional distribution of even a conditionally Gaussian ARCH process is leptokurtic. It is shown that for an ARCH(1) process, if $\alpha^2 < 1/3$ so that a finite fourth moment exists, then the kurtosis is greater than 3 for a positive α, and so the ARCH model yields observations with heavier tails than those generated by a normal distribution. Therefore, the ARCH model captures the two most common features of real high frequency financial asset returns, i.e., volatility clustering and heavy-tailed unconditional distributions.

In the ARCH*(p)* model, past shocks of more than *p* periods ago have no effect on the current volatility, hence the order *p* determines how long a shock is persistent to volatility. For financial time series, it typically requires a very high order *p* to capture the dependence. Bollerslev (1986) proposes a parsimonious way to handle with this problem, introducing the Generalised Autoregressive Conditional Heteroskedasticity (GARCH) model.

3.3.1.2 The GARCH Model

Applying the principles of the ARMA model, Bollerslev suggests a parallel proposal to the ARCH process.[5] In the GARCH(*p*,*q*) model, the conditional variance is modelled as

$$h_t = \omega + \sum_{i=1}^{p} \alpha_i \varepsilon_{t-i}^2 + \sum_{i=1}^{q} \beta_i h_{t-1} \qquad (3.11)$$

with $\omega > 0$. The GARCH process is covariance stationary if $\sum \alpha_i + \sum \beta_i < 1$. In practice, the GARCH(1,1) model is the most popular specification for estimating and forecasting volatility. The GARCH(1,1) process has just one lag of past squared error and one autoregressive term:

$$h_t = \omega + \alpha \varepsilon_{t-1}^2 + \beta h_{t-1}. \qquad (3.12)$$

[5] Taylor (1986) also proposes a similar model in an independent study.

Here α and β are non-negative to ensure the strict positivity of h_t. Note that for the GARCH(p,q) model, the positivity constraints are much more complex.[6] The conditional variance is a weighted average of three different variables: a constant ω, a forecast that was made in the previous period h_{t-1} and new information unavailable last period ε_{t-1}^2. The GARCH(1,1) process can capture a very high order of lags *p* of the ARCH(*p*) model. Indeed, by recursive substitution, the GARCH(1,1) model can be alternatively represented in the form of an ARCH(∞) process:

$$h_t = \frac{\omega}{1-\beta} + \alpha \sum_{i=1}^{\infty} \beta^{i-1} \varepsilon_{t-i}^2. \tag{3.13}$$

It obviously follows that the GARCH model is also an exponentially weighted moving average process. However, there are two major differences between the GARCH and the EWMA models. First, while the parameter λ of the EWMA process is often set *ad hoc*, the parameters of the GARCH process have to be estimated by rigorously statistical methods, normally using the Maximum Likelihood procedure. Second, the GARCH model allows the volatility process to eventually revert to its long-run level. Assume that $\alpha + \beta < 1$ so that the long-run, or unconditional variance exists $\sigma^2 = \omega(1-\alpha-\beta)^{-1}$, the *h*-step ahead forecast, by recursive substitution, is then given by

$$h_{t+h} = \sigma^2 + (\alpha+\beta)^{h-1}\left(h_{t+1} - \sigma^2\right) \tag{3.14}$$

It is inferred from (3.14) that when $\alpha + \beta < 1$, $\alpha + \beta$ dies out quickly at an exponential rate as the horizon *h* increases, hence h_{t+h} will revert to its long-run mean σ^2. Note that while α determines how fast the conditional variance responds to new information, $\alpha + \beta$ governs how fast it reverts to its long-run average. In the alternative case when $\alpha + \beta = 1$, the volatility dynamics will not converge and have to be modelled by different models.

With serially uncorrelated returns, the optimal variance forecast over the *k* cumulative steps is then given by

[6] See Nelson and Cao (1992) for more details.

$$\sum_{i=1}^{k} h_{t+i} = k\sigma^2 + \left(h_{t+1} - \sigma^2\right)\frac{1-(\alpha+\beta)^k}{1-\alpha-\beta}. \tag{3.15}$$

3.3.1.3 Other GARCH Models

The dynamics of the GARCH(1,1) model allow for intuitive interpretations as they capture some of the most important features of volatility, e.g., high persistence, fat tails and mean reversion. Moreover, it may be readily extended to richer specifications to account for other volatility features. This section gives some generalisations of the GARCH model.

The Integrated GARCH Model

The GARCH(1,1) model assumes $\alpha + \beta < 1$ so that forecasts of the conditional variance will revert to the long-run volatility level. However, it is commonly found in empirical research that volatility is so highly persistent that the sum of the estimated GARCH parameters is very close to one. Taylor (1986) estimates the GARCH(1,1) model for 40 different time series and finds that for all but six of the 40 series, the sum of the estimated parameters is equal or greater than 0.97. Engle and Bollerslev (1986) then propose the Integrated GARCH (IGARCH) to model this long-run volatility persistence.

The IGARCH(1,1) model is constructed similarly to the ARIMA model for the conditional mean, thus being considered a non-stationary GARCH(1,1) version where $\alpha + \beta = 1$. Putting $\beta = \lambda$, hence $\alpha = 1 - \lambda$, the IGARCH(1,1) process is specified by

$$h_t = \omega + (1-\lambda)\varepsilon_{t-1}^2 + \lambda h_{t-1} \tag{3.16}$$

When $\alpha + \beta = 1$, $\sigma^2 = \omega(1-\alpha-\beta)^{-1} \to \infty$, and so the IGARCH process has no finite unconditional variance. Note that the Riskmetrics EWMA model is a special case of the IGARCH(1,1) model without the drift term ω. The IGARCH(1,1) model can also be expressed as an exponentially weighted moving average process:

$$h_t = \omega + (1-\lambda)\sum_{i=1}^{\infty} \lambda^{i-1}\varepsilon_{t-i}^2 . \tag{3.17}$$

As with the EWMA model, while a shock to volatility in the IGARCH process will eventually die out at an exponential rate, it nevertheless has a permanent effect on forecast volatility at all horizons. The h-step ahead variance forecast is given by

$$h_{t+h} = h\omega + h_{t+1}. \tag{3.18}$$

Although the IGARCH model is considered a non-stationary version of the GARCH model for the conditional variance as the ARIMA model is a non-stationary version of the ARMA model for the conditional mean, there are some interesting twists in the case of the conditional variance. Nelson (1990) shows that the IGARCH(1,1) model with drift is strictly stationary, ergodic and the IGARCH(1,1) model without drift goes to zero almost surely even though it is not covariance stationary. Though the model is called Integrated GARCH, it does not follow that ε_t^2 behaves like an integrated process; on the contrary, it is still a martingale difference process. Also, while the effect of a shock is the same to both the expectation and the true process for a random walk in mean, a shock in the IGARCH process may permanently affect the expectation of a future conditional variance process, but it does not permanently affect the 'true' conditional variance process itself.

Though the IGARCH (EWMA) model may be counterfactually non-stationary, it generates better volatility forecasts than those produced by the stationary GARCH model in many empirical studies. This may be owing to the fact that IGARCH processes are not constrained by a mean level of volatility and hence can be readily adjusted to changes in unconditional volatility.

Asymmetric GARCH Models

The GARCH model suggests a symmetric volatility response to market news. The unexpected return ε_t enters the conditional variance as a square, making no difference between a positive or negative shock. However, empirical evidence suggests that in equity markets, negative shocks normally have larger effects on volatility than positive shocks of the same magnitude. Research has thus been extended to accommodate the asymmetric volatility response to market shocks, including the Exponential GARCH (EGARCH) model of Nelson (1991), the GJR-GARCH model of Glosten et al.(1993), and the Threshold GARCH (TGARCH) model of Zakoian (1994).

The Exponential GARCH (EGARCH) model

The EGARCH model of Nelson (1991) is the first asymmetric GARCH model. Instead of using the squared residuals, Nelson develops his model around the logarithmic conditional variance. The EGARCH(1,1) model also takes a different functional form:

$$\log h_t = \omega + \alpha g(z_{t-1}) + \beta \log h_{t-1} \tag{3.19}$$

where the logarithm of h_t is a function of the past $z_t \equiv \left(\sqrt{h_t}\right)^{-1} \varepsilon_t$, $g(.)$ is an asymmetric response function $g(z_t) = \theta z_t + \gamma\left[|z_t| - E|z_t|\right]$, where θ and γ are real constants. This specification enables the conditional variance to response asymmetrically to rises and falls in ε_t, since for $z_t > 0$ and $z_t < 0$, $g(z_t)$ will have different slopes, $\theta + \gamma$ and $\theta - \gamma$, respectively. The EGARCH model can capture the magnitude, as well as the sign of past shocks to volatility. Besides, by formulating conditional variance in the logarithmic form, the EGARCH model ensures that the conditional variance is positive, hence ruling out the necessity of imposing non-negative constrains on the parameters as in the GARCH model.

The conditional volatility forecast of the EGARCH process is readily available in logarithmic form. However, interests normally focus on the conditional volatility, not on the logarithmically conditional volatility. The transformation from $\log h_{t+h}$ to h_{t+h}, nevertheless, requires the entire *h*-step ahead forecast distribution of the return series. As a result, the solution is not generally available in closed form and normally derived based on rigorous procedures, such as the Monte Carlo simulation. Other models have thus been suggested to provide more straightforward specifications to forecast asymmetric conditional volatility.

The GJR-GARCH and Threshold GARCH Models

The GJR-GARCH model of Glosten et al. (1993) extends the GARCH model by still allowing quadratic response of volatility to news, but adding another ARCH term to account for asymmetric response to good and bad news. The conditional variance under the GJR-GARCH(1,1) process is specified as:

$$h_t = \omega + \alpha\varepsilon_{t-1}^2 + \delta D_{t-1}\varepsilon_{t-1}^2 + \beta h_{t-1} \tag{3.20}$$

$$\text{where } \begin{cases} D_{t-1} = 1 \text{ if } \varepsilon_t < 0 \\ D_{t-1} = 0 \text{ if } \varepsilon_t \geq 0 \end{cases} \tag{3.21}$$

To ensure the positivity of the conditional variance, ω is positive, while α, β and $\alpha+\delta$ are non-negative. It immediately follows that when $\delta>0$, negative shocks will have a higher impact on volatility than positive shocks.

The Threshold GARCH (TGARCH) model of Zakoian (1994) is constructed similarly to the GJR model, but it is formulated with the conditional standard deviation instead of the conditional variance. The TGARCH(1,1) model is given by

$$\sqrt{h_t}=\omega+\alpha\varepsilon_{t-1}+\delta D_{t-1}\varepsilon_{t-1}+\beta\sqrt{h_{t-1}}. \tag{3.22}$$

Forecasts of the GJR-GARCH and TGARCH models are straightforward to estimate. Assume further that $P\left(z_t<0\right)=P\left(z_t\geq 0\right)=0.5$, the *h*-step ahead variance forecast of the GJR process is given by

$$h_{t+h}=\sigma^2+\left(\alpha+0.5\delta+\beta\right)^{h-1}\left(h_{t+1}-\sigma^2\right) \tag{3.23}$$

with the unconditional volatility $\sigma^2=\omega\left(1-\alpha-0.5\delta-\beta\right)^{-1}$.

The GARCH model is also generalised to account for long memory behaviour in volatility. Details of the long memory volatility models are summarised in Section 3.4. Other developments of the GARCH model include, but are not restricted to, the ARCH-in-Mean model of Engle et al. (1987), the Asymmetric GARCH model of Engle and Ng (1993) and the Quadratic GARCH model of Sentana (1995).

3.3.2 Multivariate GARCH Models

As with univariate GARCH models, multivariate GARCH processes have attracted a huge interest. This section investigates some multivariate GARCH specifications, especially those that can be applied in vast dimensions. Assuming zero conditional mean, the expressions in (3.8) and (3.9) can be generalised as:

$$\mathbf{r}_t=\mathbf{H}_t^{\frac{1}{2}}\mathbf{z}_t, \tag{3.24}$$

where $\mathbf{H}_t$ is the conditional covariance matrix and $\mathbf{z}_t$ is a vector of white noise process with $E\left(\mathbf{z}_t\right)=0$ and $\operatorname{var}\left(\mathbf{z}_t\right)=\mathbf{I}_n$. Estimating the conditional covariance matrix is, inherently, challenging. The conditional covariance matrix has $\frac{1}{2}n(n+1)$ distinct parameters and structure has to be imposed to guarantee the positivity of all these parameters. Great efforts have been devoted to modelling multivariate GARCH processes, e.g. the full parameterisation VEC model of Engle and Kroner (1995), the positive definite parameterisation BEKK model (named after Baba, Engle, Kraft and Kroner), the Constant Conditional Correlation (CCC) model of Bollerslev (1990), and the Dynamic Conditional Correlation (DCC) model of Engle (2002). Section 3.3.2.1 introduces some of the common multivariate GARCH models. Due to its popular use, the DCC model is separately presented in Section 3.3.2.2.

3.3.2.1 Multivariate GARCH Models

The full parameterisation VEC representation

The full parameterisation, or VEC, representation, introduced in Engle and Kroner (1995), is the most general formulation of the multivariate GARCH models. The model converts the conditional covariance matrix into vectors of conditional variances and covariances. Under the VEC approach, the multivariate generalisation of the GARCH(1,1) model in (3.12) is defined by

$$vech\left(\mathbf{H}_t\right) = \Omega + Avech\left(\boldsymbol{\varepsilon}_{t-1}\boldsymbol{\varepsilon}_{t-1}^{'}\right) + Bvech\left(\mathbf{H}_{t-1}\right) \tag{3.25}$$

where $\boldsymbol{\varepsilon}_t = \mathbf{H}_t^{\frac{1}{2}}\mathbf{z}_t$, *vech* denotes the operator that converts the unique upper triangular elements of a symmetric matrix into a $\frac{1}{2}n(n+1)\times 1$column vector, Ω is a $\frac{1}{2}n(n+1)\times 1$ column vector, and A and B are $\frac{1}{2}n(n+1)\times\frac{1}{2}n(n+1)$ matrices. In a similar approach, we can generalise the VEC representation to the integrated or asymmetric GARCH models. Forecasts of the conditional covariance matrix can also be estimated using a recursive procedure as with the univariate models.

However, notice the number of parameters to be estimated in the full model of (3.25), which is equal to $\frac{1}{2}n^4 + n^3 + n^2 + \frac{1}{2}n = O\left(n^4\right)$. For a 25-asset conditional covariance matrix, the full model has 211,575 parameters! This is infeasible to estimate in practice. Moreover, without any additional structure imposed on the model, there is little chance that all conditional variances are positive. Therefore, several simplifications have been developed to guarantee the semi-definite positivity of the covariance matrix and to reduce the number of parameters to a manageable level.

The BEKK representation

The BEKK representation, discussed in Engle and Kroner (1995), provides a convenient way to impose restrictions on the VEC representation. The BEKK(1,1) model is given by

$$\mathbf{H}_t = \Omega + A\left(\boldsymbol{\varepsilon}_{t-1}\boldsymbol{\varepsilon}_{t-1}^{'}\right)A^{'} + B\mathbf{H}_{t-1}B^{'} \tag{3.26}$$

where Ω, *A* and *B* are symmetric positive definite $n \times n$ matrix. It is clear that the conditional covariance matrix $\mathbf{H}_t$ in (3.26) is positive definite under very weak assumptions. (3.26) is also sufficiently general when it allows all the variances and covariances to influence one another. More restrictions can be imposed in the BEKK model. In the *diagonal* BEKK model, *A* and *B* matrices are assumed to be diagonal, in which each element of the conditional covariance matrix $\mathbf{H}_t$ only depends on its own lagged values. The dynamics of variance depends only on its past variances, and the dynamics of covariance depends only on its past covariances. The BEKK representation is simplified further in the *scalar* BEKK model, where *A* and *B* matrices reduce to single values of α and β:

$$\mathbf{H}_t = \Omega + \alpha\boldsymbol{\varepsilon}_{t-1}\boldsymbol{\varepsilon}_{t-1}^{'} + \beta\mathbf{H}_{t-1}. \tag{3.27}$$

The Orthogonal GARCH model

Another way to reduce the number of estimated parameters is to impose a factor structure on the covariance matrix. Arguing that in a highly correlated system, only a few factors are required to accurately represent the system variations, Alexander (2001) proposes the Orthogonal GARCH model that combines conditional GARCH volatilities in an orthogonal Principal Component structure.

Using a Principal Components Analysis, the covariance matrix with *k* factors can be represented as

$$\mathbf{H} = \mathbf{V}\boldsymbol{\Lambda}\mathbf{V}^{'} + \mathbf{H}_{\varepsilon} \tag{3.28}$$

where **V** is an $n \times k$ matrix of factor weights, $\boldsymbol{\Lambda}$ is a diagonal matrix of the variances of the *k* factors/principal components, and $\mathbf{H}_{\varepsilon}$ is the covariance matrix of the error terms. Ignoring $\mathbf{H}_{\varepsilon}$ gives the approximation:

$$\mathbf{H} = \mathbf{V}\boldsymbol{\Lambda}\mathbf{V}^{'}. \tag{3.29}$$

As the principal components are orthogonal, the estimation of the covariance matrix reduces to the estimation of the orthogonal principal components' variances, which significantly enhances the computational efficiency. Alexander suggests that different conditional volatility EWMA or GARCH models can be employed to estimate the variances of the components. Note that the covariance matrix in (3.29) is positive semi-definite, but not strictly positive definite as there is no guarantee that $\mathbf{V\Lambda V}'$ is strictly positive definite when the number of factors is less than the number of assets.

3.3.2.2 The Dynamic Conditional Correlation Model

An alternative way to model multivariate GARCH processes in large systems is to model volatilities and correlations separately. Note that the conditional covariance matrix can be decomposed as:

$$\mathbf{H}_t = \mathbf{D}_t\mathbf{R}_t\mathbf{D}_t \tag{3.30}$$

where $\mathbf{R}_t$ is the conditional correlation matrix, $\mathbf{D}_t$ is a diagonal matrix with the standard deviations $\sqrt{h_i}$ on the i^{th} diagonal, i.e., $\mathbf{D}_t = diag\left\{\sqrt{h_i}\right\}$. In the Constant Conditional Correlation (CCC) model of Bollerslev (1990), the conditional correlation matrix $\mathbf{R}_t$ is assumed constant $\mathbf{R}_t = \mathbf{R}$ and the variations in the covariance matrix are only driven by the variations in the conditional variances. The assumptions reduce the estimation of the covariance matrix into two steps. First, a univariate GARCH model is estimated to each return series, and estimates are combined to form the diagonal matrix $\mathbf{D}_t$. Second, returns are divided by their conditional volatility to obtain the standardised, zero-mean residuals $\mathbf{e}_t = \mathbf{D}_t^{-1}\mathbf{r}_t$. The constant correlation matrix $\mathbf{R}$ is then given by the sample analogue $\mathbf{R} = T^{-1}\sum_t \mathbf{e}_t\mathbf{e}_t'$. The model is simple to estimate, and more importantly, it follows that the conditional covariance matrix $\mathbf{H}_t$ will be positive definite as long as each of the n conditional variances is well defined and the correlation matrix is positive definite.

Bollerslev suggests a convenient framework to estimate and forecast the conditional covariance matrix in large systems. However, the assumption of constant conditional correlation may be too restrictive and not suitable in many practical applications. Generalising the CCC model, Engle (2002) develops the Dynamic Conditional

Correlation (DCC) model with time-varying correlations. In the DCC model, the conditional correlation matrix is given by

$$\mathbf{R}_t = diag\left\{\mathbf{Q}_t\right\}^{-\frac{1}{2}} \mathbf{Q}_t diag\left\{\mathbf{Q}_t\right\}^{-\frac{1}{2}} \tag{3.31}$$

$$\mathbf{Q}_t = \Omega + \alpha \mathbf{e}_{t-1}\mathbf{e}_{t-1}^{'} + \beta \mathbf{Q}_{t-1}, \tag{3.32}$$

where $\mathbf{Q}_t$ is the approximation of the conditional correlation matrix $\mathbf{R}_t$. In this DCC model, $\mathbf{Q}_t$ converges to the unconditional average correlation $\overline{\mathbf{R}} = \frac{1}{T}\sum \mathbf{e}_{t-1}\mathbf{e}_{t-1}^{'}$, and $\Omega = (1-\alpha-\beta)\overline{\mathbf{R}}$. This model is an analogy to the scalar multivariate GARCH(1,1) model (see (3.27)) but in terms of volatility-adjusted returns. The positive semi-definite feature of $\mathbf{Q}_t$ is guaranteed if α and β are positive with $\alpha + \beta < 1$ and the initial matrix $\mathbf{Q}_1$ is positive definite.

Again, each conditional volatility in $\mathbf{D}_t$ can be estimated employing any univariate conditional volatility model. Returns are then divided by their conditional volatility, and the standardised, zero-mean residuals $\mathbf{e}_t = \mathbf{D}_t^{-1}\boldsymbol{r}_t$ are used to compute the quasi-conditional correlation matrix $\mathbf{Q}_t$. As the diagonal elements of $\mathbf{Q}_t$ are equal to unity only on average, $\mathbf{Q}_t$ is rescaled to obtain the conditional correlation matrix $\mathbf{R}_t = diag\left\{\mathbf{Q}_t\right\}^{-\frac{1}{2}} \mathbf{Q}_t diag\left\{\mathbf{Q}_t\right\}^{-\frac{1}{2}}$. The conditional volatility $\mathbf{D}_t$ and conditional correlations $\mathbf{R}_t$ are then combined to estimate the conditional covariance matrix $\mathbf{H}_t$.

The *h*-step-ahead conditional covariance matrix is given by

$$\mathbf{H}_{t+h} = \mathbf{D}_{t+h}\mathbf{R}_{t+h}\mathbf{D}_{t+h} \tag{3.33}$$

$\mathbf{D}_{t+h}$ is, again, estimated using the forecast procedure of the univariate volatility models. Since $\mathbf{R}_t$ is a non-linear process, the *h*-step forecast of $\mathbf{R}_t$ is not straightforward and cannot be computed using a recursive procedure. Assuming for simplicity that $E_t\left(\mathbf{e}_{t+1}\mathbf{e}_{t+1}^{'}\right) \approx \mathbf{Q}_{t+1}$, Engle and Shephard (2001) show that the forecasts of $\mathbf{Q}_{t+h}$ and $\mathbf{R}_{t+h}$ are given by

$$\mathbf{Q}_{t+h} = \sum_{j=0}^{h-2}\left(1-\alpha-\beta\right)\overline{\mathbf{Q}}\left(\alpha+\beta\right)^{j} + \left(\alpha+\beta\right)^{h-1}\mathbf{Q}_{t+1} \tag{3.34}$$

$$\mathbf{R}_{t+h} = diag\left\{\mathbf{Q}_{t+h}\right\}^{-\frac{1}{2}} \mathbf{Q}_{t+h} diag\left\{\mathbf{Q}_{t+h}\right\}^{-\frac{1}{2}}. \tag{3.35}$$

The DCC model is the most widely used multivariate GARCH model, especially in large systems, owing to its simple estimation. Specifically, we only need to estimate *n* univariate GARCH processes plus two additional parameters in (3.32). Different GARCH processes may be applied to different return series. The DCC structure is also readily flexible to allow for richer specifications. For example, in the Asymmetric DCC of Cappiello et al. (2003), an additional term is added in (3.32) so that the model allows correlation to rise more when both returns are falling than when they are both rising.

3.4 Long Memory Volatility Models

In all the conditional volatility models described above, elements of the conditional covariance matrix are typically estimated as exponentially weighted moving averages of the squares and cross products of returns. However, ample empirical evidence now suggests that although volatility is almost certainly stationary, the autocorrelation functions of the squares and cross-products of returns decline more slowly than the geometric decay rate of the EWMA and GARCH models, and hence volatility shocks are more persistent than these models imply (see, for example, Taylor, 1986, Ding et al., 1993, Andersen et al., 2001). Baillie (1996) suggests the volatility process is in a halfway house between I(0) and I(1). This empirical evidence has prompted the development of volatility models that incorporate long memory in volatility dynamics, either explicitly or implicitly. The explicit approach is to develop a model that produces hyperbolic decay in volatility's autocorrelation functions, such as the Fractionally Integrated GARCH (FIGARCH) model of Baillie et al. (1996) and the Hyperbolic GARCH (HYGARCH) model of Davidson (2004). Long memory volatility can be modelled in an implicit way, in which a combination of short memory volatility processes can generate *spurious* long memory behaviour, such as in the structural break, regime switching or component volatility models. As with the GARCH family, the Moving Average framework has been extended to allow for long memory volatility dynamics by Zumbach (2006), who develops a long memory EWMA model in which the dynamic process for volatility is defined as the logarithmically weighted sum of standard EWMA processes at different geometric time horizons. Like the short memory EWMA model of JP Morgan (1994) on which it is based, the long memory EWMA model has a highly parsimonious specification, which facilitates its implementation in practice.

This section presents some popular long memory GARCH models. The long memory EWMA model will be discussed in Chapter 5. Section 3.4.1 introduces the FIGARCH model, the most commonly used and tested long memory volatility model in the literature. Section 3.4.2 describes the Hyperbolic GARCH model that nests the GARCH, FIGARCH and IGARCH models. It allows for a more flexible dynamic structure than the FIGARCH model and facilitates tests of short versus long memory in volatility dynamics. Implicit long memory volatility models are analysed in Section 3.4.3. Section 3.4.4 gives some comments on the multivariate long memory volatility models.

3.4.1 The Fractionally Integrated GARCH Model

Baillie et al. (1996) propose the Fractionally Integrated GARCH, or FIGARCH, model, a direct conditional volatility analogy to the conditional mean ARFIMA model. In the FIGARCH model, long memory is introduced through a fractional difference operator, *d*. This model incorporates a slow hyperbolic decay for lagged squared innovations in the conditional variance while still letting the cumulative impulse response weights tend to zero, thus yielding a strictly stationary process. The conditional volatility of a FIGARCH(1,*d*,1) is given by

$$h_t = \omega + [1 - \beta L - (1 - \phi L)(1 - L)^d]\varepsilon_t^2 + \beta h_{t-1}. \tag{3.36}$$

When $d = 0$, the FIGARCH process reduces to the GARCH process. The FIGARCH model also encompasses the IGARCH model with $d = 1$. Baillie et al. (1996) show that for $0 < d \leq 1$, the FIGARCH process has no finite unconditional variance, and is not weakly stationary, the same feature with the IGARCH process. However, they show that the FIGARCH process is still strictly stationary and ergodic by a direct extension of the proof for the IGARCH case.

The one-step ahead forecast is given by

$$h_{t+1} = \omega(1 - \beta)^{-1} + [1 - (1 - \beta L)^{-1}(1 - \phi L)(1 - L)^d]\varepsilon_t^2, \tag{3.37}$$

and the *h*-step ahead volatility forecast by

$$h_{t+h} = \omega(1-\beta)^{-1} + [1-(1-\beta L)^{-1}(1-\phi L)(1-L)^{d}]\varepsilon_{t+h-1}^{2}. \tag{3.38}$$

Asymmetries are also introduced into the FIGARCH by parameterising the logarithmic conditional variance as the fractionally integrated distributed lag of past values (the Fractionally Integrated Exponential - FIEGARCH model of Bollerslev and Mikkelsen, 1996, corresponding to the EGARCH model of Nelsen, 1991), or by allowing separate influences of past positive and negative innovations as in the GJR or TGARCH model.

3.4.2 The Hyperbolic GARCH Model

Davidson (2004) notes that the FIGARCH model has non-summable autocovariances, which contradicts what we know about the actual characteristics of the volatility process. In particular, FIGARCH processes are characterised through theoretical autocorrelations decaying toward zero at a polynomial rate. This decay is so slow that the autocorrelations are not absolutely summable and, therefore, the unconditional variance is not well-defined. He then suggests the Hyperbolic GARCH (HYGARCH) model as a generalisation of the FIGARCH model (and also of the GARCH and IGARCH models). The model allows for covariance stationarity while still exhibiting hyperbolic memory. The conditional volatility of the HYGARCH model is given by

$$h_t = \omega + \left[1 - \frac{1-\phi L}{1-\beta L}\left[1+\alpha\left((1-L)^{d}-1\right)\right]\right]\varepsilon_t^2. \tag{3.39}$$

The FIGARCH and GARCH models correspond to $\alpha = 1$ and $\alpha = 0$, respectively. Also, when $d = 1$, the parameter α reduces to an autogressive root, and the process in (3.39) becomes a GARCH or an IGARCH process, depending on whether $\alpha < 1$or $\alpha = 1$. Consequently, one can test for short versus long memory in volatility dynamics by testing the hypothesis $d = 1$. When $\alpha \geq 1$, (3.39) is inherently non-stationary. On the contrary, when $0 \leq \alpha < 1$, (3.39) is covariance stationary and their cumulative impulse response weights decay towards zero at a higher rate than that implied in the FIGARCH model.

3.4.3 Component, Break and Regime Switching Volatility Models

Granger (1980) shows that the aggregation of stationary short memory AR(1) processes may result in an integrated, or a long memory process. A parallel approach applies to volatility. Long memory volatility can be modelled as a combination of different short memory volatility processes as in, for example, volatility component, structural break and regime switching models.

3.4.3.1 The Component GARCH Model

Engle and Lee (1999) introduce the Component GARCH (CGARCH) model, in which the long memory volatility process h_t is modelled as the sum of a long-run (trend) q_t and a short run (transitory) s_t volatility component, each following a GARCH-type process. The GARCH(1,1) model can be rewritten as:

$$h_t - \sigma^2 = \alpha\left(\varepsilon_{t-1}^2 - \sigma^2\right) + \beta(h_{t-1} - \sigma^2)\,. \tag{3.40}$$

The CGARCH(1,1) model allows the long-run volatility $\sigma^2 \equiv q_t$ to be time-varying and follow an autoregressive process. The CGARCH(1,1) model has the following specification:

$$h_t - q_t = \alpha\left(\varepsilon_{t-1}^2 - q_{t-1}\right) + \beta(h_{t-1} - q_{t-1}) \tag{3.41}$$

$$q_t = \omega + \rho q_{t-1} + \phi(\varepsilon_{t-1}^2 - h_{t-1})\,, \tag{3.42}$$

where $s_t = h_t - q_t$ is the transitory volatility component. The volatility innovation $\varepsilon_{t-1}^2 - h_{t-1}$ drives both the permanent and the transitory components. The long run component evolves over time following an AR process with ρ close to 1, while the short run component mean reverts to zero at a geometric rate $\alpha + \beta$. It is assumed that $0 < \alpha + \beta < \rho < 1$ so that the long run component is more persistent than the short run component.

The Component GARCH model is able to capture the high persistence of volatility dynamics and is simpler to estimate than the FIGARCH model. However, it is still computationally intensive owing to its relatively high degree of parameterisation. Engle and Lee (1999) show that the component GARCH model is in fact a constrained version of the stationary GARCH(2,2) model.

The one-step ahead forecast is given by

$$h_{t+1} = q_{t+1} + \alpha\left(\varepsilon_t^2 - q_t\right) + \beta(h_t - q_t) \tag{3.43}$$

$$q_{t+1} = \omega + \rho q_t + \phi(\varepsilon_t^2 - h_t)\,, \tag{3.44}$$

and the *h*-step ahead volatility forecast by

$$h_{t+h} = q_{t+h} + \left(\alpha + \beta\right)^{h-1}(h_t - q_t) \tag{3.45}$$

$$q_{t+h} = \frac{\omega}{1-\rho} + \rho^{h-1}\left(q_t - \frac{\omega}{1-\rho}\right). \tag{3.46}$$

3.4.3.2 Structural Break Models

Lamoureux and Lastrapes (1990) argue that the high persistence of volatility dynamics may be attributed to time-varying GARCH parameters. In particular, they allow for structural breaks in the unconditional variance of the process. They then develop a more general GARCH(1,1) model with deterministic structural breaks:

$$h_t = \omega + \delta_1 D_{1t} + \ldots + \delta_k D_{kt} + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 h_{t-1}, \tag{3.47}$$

where $D_{it}\left(i = 1,\ldots,k\right)$ are dummy variables that correspond to periods over which the GARCH process is stationary. Note that there are $k+1$ such periods in (3.47). The difficulty here lies in the determination of the timing of the breaks in the unconditional variance. The one-step ahead forecast is given by

$$h_{t+1} = \omega + \delta_1 D_{1t} + \ldots + \delta_k D_{kt} + \alpha_1 \varepsilon_t^2 + \beta_1 h_t\,, \tag{3.48}$$

and the *h*-step ahead volatility forecast by

$$h_{t+h} = \omega + \delta_1 D_{1t} + \ldots + \delta_k D_{kt} + (\alpha_1 + \beta_1) h_{t+h-1}\,. \tag{3.49}$$

3.4.3.3 Regime Switching Models

An alternative to modelling long memory volatility is to use regime switching models. Hamilton and Susmel (1994) note that financial markets react to large and small shocks differently and the rate of mean reversion is faster for large shocks. They originate a new class of regime switching models, where the GARCH volatility process can take different parameter values, depending on whether it is in a high or a low volatility regime. The most general regime switching model takes the form suggested in Gray (1996):

$$\sigma^2_{t,S_{t-1}} = \omega_{S_{t-1}} + \alpha_{S_{t-1}} \varepsilon^2_{t-1} + \beta_{S_{t-1}} \sigma^2_{t-1,S_{t-1}} \tag{3.50}$$

where S_t defines the regime at time *t*. Numerous regime switching models have been, since then, developed to allow different switching probability.

3.4.4 Multivariate Long Memory Volatility Models

Though correctly specified to capture the high persistence property of volatility dynamics, long memory volatility models are often problematic to implement in practice, not least because of their complexity in estimation. For example, the FIGARCH model requires very long periods of historical data in order to calibrate the hyperbolic decay functions on which it is based. Consequently, their use is limited in many practical situations, especially where volatility forecasts are required in real time (such as on an options trading desk) or where the model must be estimated a large number of times over a rolling window (such in the back testing of risk management systems). In the multivariate context, long memory volatility modelling poses even more significant computational challenges, especially so for the high dimensional covariance matrices that are typically encountered in asset allocation and risk management. As with the GARCH models, the univariate long memory conditional volatility models can be generalised to the multivariate case using the VEC, BEKK or DCC models. However, the complexity and computational intensity have limited the use of multivariate long memory volatility models to very low dimensional systems, even though many applications in finance require forecasts of high dimensional covariance matrices. For example, Teyssiere (1998) estimates the covariance matrix for three foreign exchange return series using both an unrestricted multivariate FIGARCH model and a FIGARCH model implemented with the Constant Conditional Correlation (CCC) structure of Bollerslev (1990). Similarly, Niguez and Rubia (2006) model the covariance matrix of five foreign exchange series using an Orthogonal HYGARCH model, which combines the univariate HYGARCH long memory volatility model of Davidson (2004) with the multivariate Orthogonal GARCH framework of Alexander (2001). Zumbach (2009) develops a multivariate version of the univariate long memory EWMA model, in which elements of the covariance matrix are estimated as the averages of the squares and cross products of past returns with predetermined logarithmically decaying weights. The parsimony of the long memory EWMA model promises potentially beneficial application in high dimensional systems

3.5 Conditional Volatility Models and Asset Allocation

The literature suggests that there are significant economic benefits to exploiting the forecasts of multivariate conditional volatility models relative to using the unconditional covariance matrix in the asset allocation framework. Fleming et al. (2001, 2003) are among the first to study the economic value of exploiting conditional covariance matrices for investors. In a volatility timing asset allocation framework where investors assume constant expected returns and rebalance their portfolios based on forecasts of the covariance matrix, they show that investors are better off in terms of utility when switching from a static to a dynamic asset allocation strategy. Recent studies incorporate more properties of volatility dynamics in application to investment decisions. Thorpe and Milunovich (2007) allow for asymmetries in modelling volatilities and correlations, and show that investors are willing to pay to switch from symmetric to asymmetric forecasts. Similarly, Hyde et al. (2010) demonstrate the benefits of accounting for volatility jumps in asset allocation strategies. Conditional volatility models have also been embedded with a factor structure to reduce estimation error. For example, Briner and Connor (2008) allow for the dynamic variations of returns' volatility and covariance in a traditional factor model by imposing an exponential weighting on the factor covariance matrix. Han (2006) develops a dynamic factor multivariate stochastic volatility model, which utilises unobserved factors to capture the dynamic behaviour of volatility (and also returns) in an asset allocation problem. The research generally favours the use of the dynamic factor-structured covariance matrix to the unstructured alternatives. Owing to the complexity in estimation, long memory conditional volatility models have rarely been used in the asset allocation framework where forecasts of the high dimensional covariance matrix are normally required. The next chapters will fill in this gap, studying the benefits of allowing for long memory volatility dynamics in forecasts of the covariance matrix for asset allocation.

Chapter 4

Data Analysis

4.1 Data Description

The research first evaluates the forecast performance of a range of multivariate long memory conditional volatility models using the asset allocation framework of Engle and Colacito (2006). Details of these multivariate long memory volatility models will be provided in Chapter 5. The empirical research hence employs the same three sets of assets as in Engle and Colacito (2006). These comprise a high correlation bivariate system (the S&P500 and DJIA indices), a low correlation bivariate system (the S&P500 and 10-year Treasury bond futures), and a moderate correlation high dimensional system (21 stock international stock indices and 13 international bond indices). I additionally consider another high dimensional system, comprising the components of the DJIA index. The four datasets are also used to study the economic benefits of allowing for long memory volatility dynamics in estimating and forecasting the covariance matrix for dynamic asset allocation in Chapter 6.

The two bivariate systems are now described in detail. The low correlation Stock-Bond system uses daily data for the S&P500 index and 10-year Treasury bond futures, while the high correlation S&P500-DJIA system uses daily data for the S&P500 and Dow Jones Industrial Average indices. All data are from Datastream and cover the period 01 January 1988 to 31 December 2009. The futures prices are continuous series of futures settlement prices, starting at the nearest contract month, which forms the first values for the continuous series until either the contract reaches its expiration date or until the first business day of the notional contract month, whichever is sooner. At this point prices from the next trading contract month are taken. There may be a non-synchronicity issue in the Stock-Bond futures as the Bond and Stock futures contracts close at 2:00 CST and 3:15 CST, respectively. Returns are calculated as the log price difference over consecutive days. I exclude from the sample all days on which any of the markets was closed, yielding 5548 observations for each dataset. As the futures contracts require no initial investment, the futures returns are approximately equivalent to excess spot returns. The returns of the S&P500 and DJIA indices are converted to excess returns by subtracting the daily 1-month T-Bill rate.[7] Table 4.1 reports some descriptive statistics of the four return series. The annualised average return on Stock is nearly five times as much as the return on Bond. Expectedly, the higher return on Stock is accompanied by a higher level of risk, 19.06% as compared to that of 6.53% of Bond. The sample correlation of the stock index futures and the bond futures is very close to zero, while for the S&P500 and DJIA indices, it is close to one. As a result, the return and risk properties of the S&P500 and DJIA indices are similar, though the DJIA index performs slightly better with a higher return and lower risk. For all series, returns are negatively skewed and leptokurtic.

Following Engle and Colacito (2006), I also consider a moderate correlation high dimensional system. An international stock and bond portfolio is constructed from 34 assets, comprising 21 stock indices from the FTSE All-World indices and 13 five-year average maturity bond indices. The 21 stock indices and 13 bond indices include all of the major world stock and government bond markets. All data are taken from Datastream and converted to US dollar denominated prices. Following Engle and Colacito (2006), I use weekly returns to avoid the problem of non-synchronous trading. Weekly returns are calculated as the log price difference using Friday to Friday closing prices. The dataset comprises 22 years of weekly returns, yielding a total of 1147 observations from 01 January 1988 to 31 December 2009. Descriptive statistics for the international dataset are given in Table 4.2. For all countries for which both stock and bond indices are present, the stock index has a higher return and higher risk than the corresponding bond index. The US is the least risky market for both stocks and bonds. Smaller countries, such as Austria, Hong Kong, Ireland and Mexico generally have higher risk, although this is not always accompanied by higher returns. Japan and New Zealand have negative annualised average stock returns over the sample considered. Returns are, again, leptokurtic and, in most cases, negatively skewed. The international stock markets are relatively highly correlated, as are the international bond markets. The average correlation coefficient among the 21 stock market return series is 0.54, while among the bond market return series it is 0.61. However, the stock and bond markets as a whole have an average correlation coefficient of only 0.20.

[7] This is the simple daily rate that, over the number of trading days in the month, compounds to 1-month T-Bill rate from Ibbotson and Associates, Inc.

I additionally consider a higher frequency high dimensional system, comprising the components of the Dow Jones Industrial Average (DJIA) index as of 31 December 2009. Daily data are collected from the Centre for Research in Security Prices from 01 March 1990 to 31 December 2009. I exclude Kraft, which was listed only in June 2001. Returns are calculated as the log price difference over consecutive days. All days on which the market was closed are excluded from the sample, yielding 5001 observations. Table 4.3 provides the summary analysis of the 29 DJIA stocks. Annualised average returns are positive for all 29 stocks, with Bank of America (BAC) being the lowest (1.51%) and Cisco (C) being the highest (28.74%). The return series are again highly non-normal, with very high leptokurtosis. The average correlation coefficient of the DJIA components is 0.34.

4.2 Evidence of Long Memory in Volatility

Figure 4.1 plots the sample autocorrelations of returns, absolute returns and squared returns for the four return series of Stock, Bond, S&P500 and DJIA. While the autocorrelations of normal returns are not significantly different from zero, the autocorrelations of absolute returns and squared returns are highly persistent and still positively significant up to lag 100. The autocorrelations of absolute returns are also consistently higher than those of squared returns, a feature first identified by Taylor (1986). The slowly decaying autocorrelation functions of absolute returns and squared returns suggest the presence of long memory in volatility.

Formal tests are conducted to confirm the visual evidence of long memory in volatility, the results of which are reported in Table 4.4.[8] The parametric FIGARCH model is estimated for the whole sample, and the estimated fractional difference orders $I(d)$ range from 0.35 to 0.49. Semi-parametric long memory tests such as the narrow band log periodogram (GPH) estimator of Geweke and Porter-Hudak (1983) and the broad band log periodogram (MS) estimator of Moulines and Soulier (1999) are also applied. To estimate the GPH and MS operators, I use the recommended bandwidth m equal to the square root of the sample size $(m = 77)$ and the Fourier term p equal to the log of the sample size $(p = 4)$, respectively. The table reports the results for both squared returns and absolute returns. All the tests suggest long memory in volatility for all four series and that stock return volatility has longer memory than bond return volatility.

[8] The tests are conducted using the Time Series Modelling software of James Davidson.

Consistent with the graphical results of the autocorrelations, absolute returns demonstrate a higher level of persistence than squared returns. The one-sided tests for the hypothesis $d = 0.5$ are conducted against the alternative $d < 0.5$. Rejecting the hypothesis, I confirm that the volatility processes of all four series are characterised by long memory, but are nevertheless stationary.

The presence of long memory volatility is now examined in the multivariate systems. Table 4.5 reports the sum of the first 100 autocorrelation coefficients of squared returns and of absolute returns for some return series and their average values for the 21 international stock indices, the 13 international bond indices, and the 29 stocks of the DJIA index. The table also reports the fractional difference operators estimated using the FIGARCH, the GPH and the MS tests. To conduct the GPH and MS tests, I use the recommended bandwidths $m = 33$ and the Fourier terms $p = 3$ for the international stock and bond portfolio, and the corresponding values of $m = 71$, and $p = 4$ for the DJIA portfolio. All return series show long memory behaviour in volatility. Again, the level of persistence of absolute returns is consistently higher than that of squared returns, which is clearly demonstrated in both the sum of autocorrelation coefficients and the fractional difference operators. For all countries for which both stock and bond indices are present, stock index volatility is also more persistent than the corresponding bond index volatility. The average fractional difference operator of squared returns on the stock indices is 0.44 with the parametric FIGARCH test, and 0.32 with the semi-parametric GPH tests, while for the international bond indices, the corresponding results are 0.30 and 0.25. Long memory volatility is also clearly evident in the individual DJIA stocks, with the average fractional difference orders of 0.37 with the FIGARCH test, and of 0.42 with GPH test. Based on the standard errors, not reported here, all the fractional difference operators estimated using both the GPH and MS tests are significantly greater than zero.

Table 4.1. Summary Statistics for the Two Bivariate Systems

The table reports descriptive statistics for the daily returns on Stock and Bond futures, and the daily excess returns on the S&P500 and DJIA indices. Means and standard deviations are annualised. The sample period is from 01 January 1988 to 31 December 2009. The table also reports the statistics for the Jarque-Bera tests of the null hypothesis that the series follows normal distribution. All the statistics confirm the rejection of the normality hypothesis at 1% significance level.

Return series	Mean (%)	Std. Dev. (%)	Skewness	Kurtosis	Min (%)	Max (%)	Normality test	Corr.
Stock	6.83	19.06	-0.19	14.18	-10.40	13.20	28936	-0.04
Bond	1.48	6.53	-0.28	6.63	-2.86	3.57	3123	
S&P500	2.80	18.34	-0.25	12.32	-9.47	10.95	20117	0.96
DJIA	3.59	17.72	-0.20	11.62	-8.20	10.51	17194	

Table 4.2. Summary Statistics for the International Stock and Bond Returns

The table reports summary statistics for the weekly returns on 21 international stock indices and 13 government bond indices. Means and standard deviations are annualised. The sample period is from 01 January 1988 to 31 December 2009. The table also reports the statistics for the Jarque-Bera tests of the null hypothesis that the series follows normal distribution. All the statistics, except for those with *, confirm the rejection of the normality hypothesis at 1% significance level. * denotes rejection of the normality hypothesis at 5% significance level.

Return series	Mean (%)	Std. Dev. (%)	Skewness	Kurtosis	Min (%)	Max (%)	Normality test
Panel A. International Stocks							
Australia	7.30	21.64	-1.77	21.33	-34.86	14.52	16657
Austria	6.03	25.81	-1.52	18.70	-38.22	20.94	12223
Belgium	5.45	20.98	-1.21	12.68	-26.88	12.53	4757
Canada	7.58	20.68	-1.13	13.91	-25.92	17.61	5930
Denmark	9.89	21.00	-1.31	13.35	-26.39	13.66	5446
France	7.44	21.25	-0.90	10.94	-27.16	13.76	3167
Germany	6.69	23.49	-0.80	8.93	-26.11	15.00	1800
Hongkong	9.22	25.37	-0.62	6.57	-21.08	13.85	682
Ireland	3.44	25.49	-1.72	19.88	-39.31	16.18	14184
Italy	2.91	24.82	-0.60	8.85	-26.71	19.04	1705
Japan	-1.29	22.56	0.07	4.67	-16.02	11.75	134
Mexico	19.21	33.81	-0.33	7.66	-30.20	23.23	1060
Netherland	6.88	20.89	-1.44	17.48	-31.48	14.85	10416
New Zealand	-0.08	22.04	-0.63	7.44	-23.06	12.07	1017
Norway	8.93	26.82	-0.84	10.37	-28.54	19.82	2733
Singapore	6.99	26.29	-0.69	13.21	-33.13	23.02	5071
Spain	6.92	22.28	-0.90	10.21	-26.22	13.76	2641
Sweden	9.79	26.88	-0.52	7.73	-25.12	19.05	1123
Switzerland	8.44	19.42	-0.70	11.14	-24.01	13.96	3263
UK	4.44	19.13	-1.05	16.81	-27.73	16.30	9324
US	7.03	16.81	-0.81	10.54	-20.19	11.45	2845
Panel B. International Bonds							
Austria	0.92	10.58	-0.03	3.64	-5.85	5.72	20
Belgium	0.95	10.68	-0.02	3.47	-5.16	5.55	11
Canada	2.36	8.71	-0.51	6.53	-8.38	5.34	647
Denmark	1.60	10.92	0.00	3.84	-5.82	5.67	33
France	1.81	10.54	-0.02	3.47	-4.88	5.79	11
Germany	0.73	10.62	0.01	3.37	-4.52	5.77	7*
Ireland	1.83	10.89	-0.25	4.19	-7.52	5.94	79
Japan	1.67	12.11	0.89	8.33	-6.05	14.30	1509
Netherland	0.55	10.64	-0.02	3.36	-4.82	5.45	6*
Sweden	0.06	12.06	-0.18	3.84	-7.85	5.93	40
Switzerland	0.95	12.05	0.11	3.72	-6.28	6.89	27
UK	0.13	10.60	-0.24	4.93	-7.12	6.48	188
US	1.23	4.43	-0.19	3.82	-2.61	2.06	39

Table 4.3. Summary Statistics for the DJIA Components

The table reports summary statistics for the daily returns on the 29 components of the DJIA index. Means and standard deviations are annualised. The sample period is from 01 March 1990 to 31 December 2009. The table also reports the statistics for the Jarque-Bera tests of the null hypothesis that the series follows normal distribution. All the statistics confirm the rejection of the normality hypothesis at 1% significant level.

Return series	*Mean (%)*	*Std. Dev. (%)*	*Skewness*	*Kurtosis*	*Min (%)*	*Max (%)*	*Normality test*
AA	3.47	39.12	-0.02	11.23	-17.50	20.87	14102
AXP	7.22	38.76	0.03	9.94	-19.35	18.77	10043
BA	4.61	31.89	-0.33	9.73	-19.39	14.38	9525
BAC	1.51	45.21	-0.29	30.90	-34.21	30.21	162245
CAT	10.19	33.62	-0.08	7.18	-15.69	13.74	3652
C	28.74	46.95	0.00	7.48	-22.10	21.82	4175
CVX	7.62	25.60	0.13	12.63	-13.34	18.94	19331
DD	2.76	29.39	-0.09	7.10	-12.03	10.86	3513
DIS	6.43	32.11	0.00	10.40	-20.29	14.82	11410
GE	5.44	29.93	0.01	11.17	-13.68	17.98	13916
GM	13.54	35.10	-0.67	16.81	-33.88	13.16	40119
HD	11.51	40.37	-0.08	9.21	-20.70	18.99	8044
HPQ	8.14	30.53	0.04	9.76	-16.89	12.37	9537
IBM	13.99	42.72	-0.38	8.26	-24.89	18.33	5884
INTC	11.37	23.70	-0.19	9.75	-17.25	11.54	9510
JNJ	8.01	42.10	0.26	13.11	-23.23	22.39	21336
JPM	9.39	24.79	0.08	8.01	-11.07	13.00	5230
KO	10.43	26.89	-0.04	6.98	-13.72	10.31	3305
MCD	7.12	24.25	0.01	7.50	-10.08	10.50	4214
MMM	5.82	29.95	-1.09	22.53	-31.17	12.25	80485
MRK	19.06	35.23	0.01	7.94	-16.96	17.87	5087
MSFT	10.01	29.62	-0.18	6.07	-11.82	9.69	1997
PFE	10.26	25.33	-2.78	68.38	-37.66	9.73	897033
PG	3.50	28.68	0.08	7.39	-13.54	15.08	4027
T	6.01	30.34	0.34	16.22	-20.07	22.76	36490
UTX	11.97	28.77	-1.13	28.55	-33.20	12.79	137065
VZ	1.90	27.61	0.17	7.64	-12.61	13.66	4503
WMT	11.39	29.21	0.13	5.83	-10.26	10.50	1681
XOM	8.91	24.83	0.09	11.92	-15.03	15.86	16591

Table 4.4. Fractional Difference Operators for the Two Bivariate Systems

The table reports the fractional difference operators, *d*, estimated using the FIGARCH, Geweke-Porter-Hudak (GPH) and Moulines-Soulier (MS) tests. The GPH and MS estimators are applied for both squared returns and absolute returns. The standard errors are reported in parentheses.

Return series	$\hat{d}_{FIGARCH}$	**Squared returns**		**Absolute returns**	
		$\hat{d}_{GPH}$	$\hat{d}_{MS}$	$\hat{d}_{GPH}$	$\hat{d}_{MS}$
Stock	0.403	0.357	0.373	0.553	0.406
		(0.080)	(0.032)	(0.080)	(0.032)
Bond	0.355	0.410	0.190	0.456	0.209
		(0.080)	(0.032)	(0.080)	(0.032)
S&P500	0.492	0.441	0.461	0.607	0.443
		(0.080)	(0.032)	(0.080)	(0.032)
DJIA	0.487	0.396	0.417	0.584	0.427
		(0.080)	(0.032)	(0.080)	(0.032)

Table 4.5. Autocorrelations and Fractional Difference Operators for the Multivariate Systems

The table reports the sum of the first 100 autocorrelation coefficients of squared returns $\sum \rho\left(r^2\right)$ and of absolute returns $\sum \rho\left(|r|\right)$ for some return series and their average values for the 21 international stock indices, the 13 international bond indices, and the 29 stocks of the DJIA index. The fractional difference operators, *d*, are estimated using the FIGARCH, the Geweke-Porter-Hudak (GPH) and the Moulines-Soulier (MS) tests. The GPH and MS estimators are applied for both squared and absolute returns.

Return series	No. of obs	$\sum \rho\left(r^2\right)$	$\sum \rho\left(\lvert r\rvert\right)$	$\hat{d}_{FIGARCH}$	Squared returns		Absolute returns	
					$\hat{d}_{GPH}$	$\hat{d}_{MS}$	$\hat{d}_{GPH}$	$\hat{d}_{MS}$
Panel A. International Stock and Bond Portfolio								
Stock Market Indices								
France	1147	3.47	6.03	0.41	0.30	0.36	0.40	0.37
Germany	1147	3.41	5.53	0.45	0.26	0.48	0.36	0.45
Japan	1147	1.44	1.74	0.16	0.29	0.20	0.32	0.26
UK	1147	3.52	7.00	0.38	0.27	0.54	0.46	0.45
US	1147	4.02	8.92	0.42	0.37	0.34	0.55	0.37
Averages of 21 stock indices		*3.37*	*6.44*	*0.44*	*0.32*	*0.33*	*0.44*	*0.36*
Bond Indices								
France	1147	1.84	3.15	0.26	0.30	0.24	0.38	0.19
Germany	1147	1.42	2.93	0.28	0.15	0.18	0.27	0.19
Japan	1147	1.00	1.86	0.30	0.29	0.13	0.36	0.19
UK	1147	2.96	3.41	0.19	0.43	0.23	0.37	0.26
US	1147	1.69	1.42	0.32	0.17	0.15	0.06	0.13
Averages of 13 bond indices		*1.89*	*2.79*	*0.30*	*0.25*	*0.22*	*0.29*	*0.20*
Panel B. DJIA Portfolio								
AA	5001	16.56	19.21	0.33	0.43	0.35	0.54	0.34
BAC	5001	18.12	31.34	0.45	0.78	0.33	0.77	0.42
C	5001	7.52	13.11	0.42	0.59	0.21	0.67	0.30
DIS	5001	5.64	12.03	0.28	0.47	0.31	0.56	0.33
MSFT	5001	5.76	9.39	0.33	0.44	0.21	0.47	0.26
T	5001	8.81	15.42	0.37	0.49	0.35	0.63	0.38
Averages of 29 stocks		*7.80*	*13.22*	*0.37*	*0.42*	*0.25*	*0.55*	*0.30*

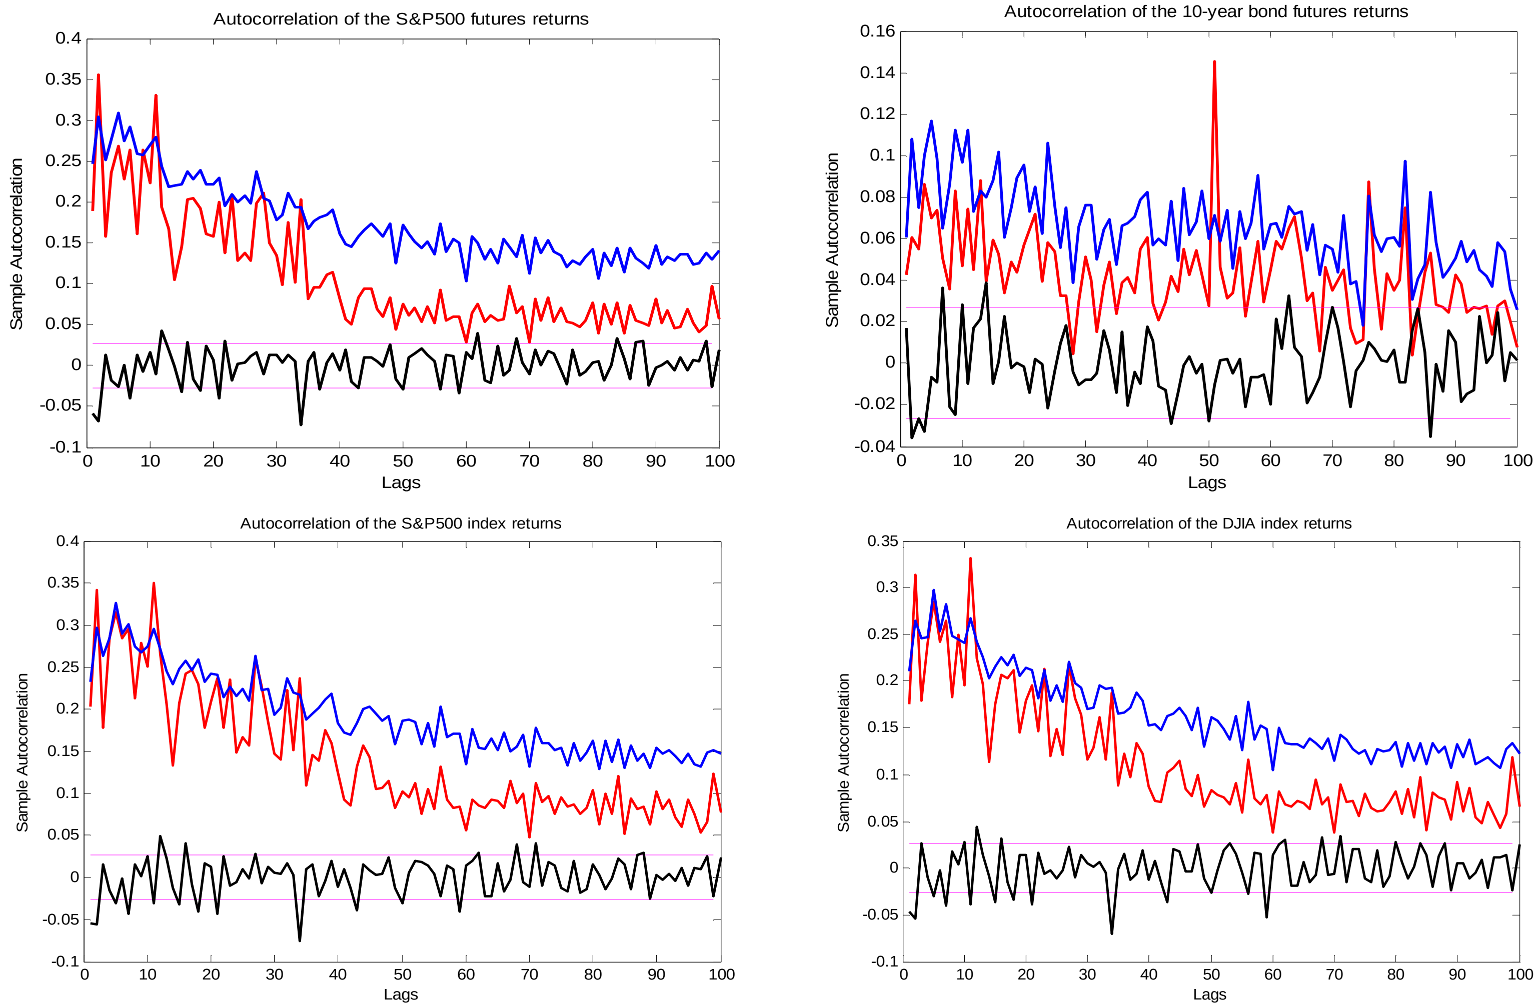

Figure 4.1. Autocorrelation of Returns (Black Line), Absolute Returns (Blue Line), and Squared Returns (Red Line)

Chapter 5

Long Memory Conditional Volatility and Asset Allocation

In this chapter, I evaluate the forecast performance of the long memory covariance matrix over both short and long horizons, using the asset allocation framework of Engle and Colacito (2006). In so doing, I compare the performance of a number of long memory and short memory multivariate volatility models. While many alternative volatility models have been developed in the literature, my choice reflects the need for parsimonious models that can be used to forecast high dimensional covariance matrices. I employ four long memory volatility models: the multivariate long memory EWMA model of Zumbach (2009b), and three multivariate long memory implemented using the Dynamic Conditional Correlation (DCC) framework of Engle (2002). These are the univariate long memory univariate EWMA model of Zumbach (2006), the component GARCH model of Engle and Lee (1999) and the FIGARCH model of Baillie et al. (1996). I compare the four multivariate long memory models with two multivariate short memory models. These are the very widely used RiskMetrics EWMA model of JP Morgan (1994), and the DCC model implemented with the univariate GARCH model.

I use the six multivariate volatility models to forecast the covariance matrices for the four datasets described in Chapter 4. These comprise low/high dimensional, low/high correlation systems. In particular, the two bivariate systems include a low correlation S&P500 and 10-year Treasury bond futures (Stock-Bond) portfolio and a high correlation S&P500 and DJIA index (S&P500-DJIA) portfolio, while the two moderate correlation, high dimensional systems consist of an international stock and bond portfolio and a US all-stock portfolio. The analysis is conducted using data over the period from 1 January 1988 to 31 December 2009, and considers forecast horizons of up to three months. For the two bivariate systems, I first evaluate the forecast performance of the models using a range of statistical criteria that measure the accuracy, bias and informational content of the models' forecasts over varying time horizons. For all four systems, I then employ Engle and Colacito's (2006) approach to assess the economic value of the forecast covariance matrices in an asset allocation setting. I report two main findings. The first is that for longer horizon forecasts, multivariate long memory volatility models generally produce forecasts of the covariance matrix that are both statistically more accurate and informative, and economically more useful than those produced by short memory volatility models. The second is that the two long memory models that are based on the Zumbach (2006) univariate model outperform the other models – both short memory and long memory – in a majority of cases across all forecast horizons. These results apply to all four datasets and are robust to the choice of estimation window.

The remainder of this chapter is organised as follows. Section 5.1 provides details of the multivariate conditional volatility models used in the empirical analysis. Section 5.2 describes the methods applied to evaluate forecast performance for the six models. In Section 5.3, I report the empirical results of the analysis, while Section 5.4 offers some concluding comments and some suggestions for future research.

5.1 Multivariate Long Memory Conditional Volatility Models

Motivated by the need for parsimonious models that can be used to forecast high dimensional covariance matrices, I first consider two simple multivariate long memory conditional volatility models based on the univariate long memory volatility model of Zumbach (2006). The first is the multivariate long memory EWMA (LM-EWMA) model of Zumbach (2009b), which is a simple multivariate extension of the univariate long memory EWMA model in which both the variances and covariances are governed by the same long memory process, and is thus the long memory analogue of the short memory multivariate RiskMetrics EWMA model of JP Morgan (1994). In the second, I employ the Dynamic Conditional Correlation framework of Engle (2002) to model the dynamic processes of the correlations directly, using the univariate long memory EWMA model for the individual variances. This is the long memory EWMA-DCC (LM-EWMA-DCC) model. I compare the two long memory EWMA models with the multivariate FIGARCH(1,d,1) and Component GARCH(1,1) (CGARCH) long memory models, both implemented using the DCC framework. To evaluate the relative benefits of allowing for long memory in forecasting the covariance matrix, I compare the four multivariate long memory volatility models with two multivariate short memory volatility models. These are the multivariate RiskMetrics EWMA model of JP Morgan (1994) and the GARCH(1,1) model implemented using the DCC framework. In this section, I give details of each of these six models.

5.1.1 The Multivariate LM-EWMA Model

Consider an n-dimensional vector of returns $\boldsymbol{r}_t = \left(r_{1t}, r_{2t}, \ldots, r_{nt}\right)'$ with conditional mean zero and conditional covariance matrix $\mathbf{H}_t$:

$$\boldsymbol{r}_t = \mathbf{H}_t^{\frac{1}{2}} \mathbf{z}_t , \tag{5.1}$$

where $\mathbf{z}_t$ is i.i.d with $E\left(\mathbf{z}_t\right) = 0$ and $\operatorname{var}\left(\mathbf{z}_t\right) = \mathbf{I}_n$. Zumbach (2009b) considers the class of conditional covariance matrices that are the weighted sum of the cross products of past returns:

$$\mathbf{H}_{t+1} = \sum_{i=0}^{\infty} \lambda(i) \boldsymbol{r}_{t-i} \boldsymbol{r}_{t-i}^{'} , \tag{5.2}$$

with $\sum \lambda(i) = 1$. In the RiskMetrics EWMA model of JP Morgan (1994), the weights $\lambda(i)$ decay geometrically, yielding a short memory process for the elements of the covariance matrix. The long memory conditional covariance matrix is defined as the weighted average of K standard (short memory) multivariate EWMA processes:

$$\mathbf{H}_{t+1} = \sum_{k=1}^{K} w_k \mathbf{H}_{k,t} \tag{5.3}$$

where

$$\mathbf{H}_{k,t} = \mu_k \mathbf{H}_{k,t-1} + \left(1 - \mu_k\right) \boldsymbol{r}_t \boldsymbol{r}_t ' . \tag{5.4}$$

The decay factor μ_k of the k^{th} EWMA process is defined by a characteristic time τ_k such that $\mu_k = \exp\left(\frac{-1}{\tau_k}\right)$, with geometric time structure $\tau_k = \tau_1 \rho^{k-1}$ for $k = (1,\ldots,K)$. Zumbach (2006) sets ρ to the value of $\sqrt{2}$. The memory of the volatility process is determined by the weights w_k, which are assumed to decay logarithmically:

$$w_k = \frac{1}{C}\left(1 - \frac{ln\left(\tau_k\right)}{ln\left(\tau_0\right)}\right) \tag{5.5}$$

with the normalization constant $C = K - \sum_k \frac{ln(\tau_k)}{ln(\tau_0)}$ such that $\sum_k w_k = 1$. The conditional covariance matrix is therefore parsimoniously defined as a process with just three parameters: τ_1 (the shortest time scale at which volatility is measured, i.e. the lower cut-off), τ_K (the upper cut-off, which increases exponentially with the number of components K), and τ_0 (the logarithmic decay factor). For the univariate case, Zumbach (2006) sets the optimal parameter values at $\tau_0 =$ 1560 days = 6 years, $\tau_1 =$ 4 days, and $\tau_K =$ 512 days, which is equivalent to $K = 15$.

The EWMA process in (5.4) can also be expressed as

$$\mathbf{H}_{k,t} = \left(1-\mu_k\right)\sum_{i=0}^{\infty}\mu_k^i \boldsymbol{r}_{t-i}\boldsymbol{r}_{t-i}^{'}. \tag{5.6}$$

Hence the LM-EWMA model can be written in the form of (5.2):

$$\mathbf{H}_{t+1} = \sum_{i=0}^{\infty}\sum_{k=1}^{K} w_k \left(1-\mu_k\right)\mu_k^i \boldsymbol{r}_{t-i}\boldsymbol{r}_{t-i}^{'} = \sum_{i=0}^{\infty}\lambda(i)\boldsymbol{r}_{t-i}\boldsymbol{r}_{t-i}^{'} \tag{5.7}$$

with $\lambda(i) = \sum_k w_k \left(1-\mu_k\right)\mu_k^i$ and $\sum_i \lambda\left(i\right) = 1$ (which is satisfied by $\sum_k w_k = 1$). When $K = 1$, the LM-EWMA process reduces to the short memory RiskMetrics EWMA process. Note that since $\mathbf{H}_{k,t}$ is a positive definite matrix (see Riskmetrics, 1994), $\mathbf{H}_{t+1}$, which is a linear combination of $\mathbf{H}_{t,k}$ with positive weights, will also be positive definite. Since the LM-EWMA covariance matrix is the sum of EWMA processes over increasing time horizons, forecasts of the covariance matrix are straightforward to obtain using a recursive procedure, which is detailed in Appendix 5.1. The one-step-ahead forecast of the covariance matrix is already given by (5.7). Under the assumption of serially uncorrelated returns, the *h*-step cumulative forecast of the covariance matrix given the information set $\mathcal{F}_t$ at time *t* is equal to:

$$\mathbf{H}_{t+1:t+h} = h\sum_{i=0}^{T}\lambda(h,i)\boldsymbol{r}_{t-i}\boldsymbol{r}_{t-i}^{'} \tag{5.8}$$

with the weights $\lambda\left(h,i\right)$ given by

$$\lambda\left(h,i\right) = \sum_{k=1}^{K}\frac{1}{h}\sum_{j=1}^{h-1} w_{j,k}\frac{\left(1-\mu_k\right)}{1-\mu_k^T}\mu_k^i \tag{5.9}$$

where T is the cut-off time, $w_{j,k}$ is the k element of vector $\boldsymbol{w}_j = \boldsymbol{w}'\left[\boldsymbol{M} + \left(\boldsymbol{\iota} - \boldsymbol{\mu}\right)\boldsymbol{w}'\right]^j$, $\boldsymbol{\mu}$ is the vector of μ_k, $\boldsymbol{M}$ is the diagonal matrix consisting of μ_k, and $\boldsymbol{\iota}$ is the unit vector. Since $\sum_k w_k = 1$, we obtain $\sum \lambda(h,i) = 1$. Also note that when $K = 1$, then $w = 1$, and so the LM-EWMA process reduces to a standard short memory EWMA process with forecast weights $\lambda\left(h,i\right) = \left(1-\mu_k\right)\mu_k^i \big/ \left(1-\mu_k^T\right)$, independent of the forecast horizon. As the weights $\lambda\left(h,i\right)$ can be estimated a priori, without reference to the data, the forecast in (5.8) is straightforward to compute. As with the standard EWMA model, the LM-EWMA model circumvents the computational burden of other multivariate long memory models, and indeed can easily be implemented in a spreadsheet.

5.1.2 The Multivariate LM-EWMA-DCC Model

In the Dynamic Conditional Correlation (DCC) model of Engle (2002), the conditional covariance matrix is decomposed as follows:

$$\mathbf{H}_t = \mathbf{D}_t \mathbf{R}_t \mathbf{D}_t \tag{5.10}$$

$$\mathbf{R}_t = diag\left\{\mathbf{Q}_t\right\}^{-\frac{1}{2}} \mathbf{Q}_t diag\left\{\mathbf{Q}_t\right\}^{-\frac{1}{2}} \tag{5.11}$$

$$\mathbf{Q}_t = \Omega + \alpha \boldsymbol{\varepsilon}_{t-1}\boldsymbol{\varepsilon}_{t-1}' + \beta \mathbf{Q}_{t-1} \tag{5.12}$$

where $\mathbf{R}_t$ is the conditional correlation matrix, $\mathbf{D}_t$ is a diagonal matrix with the time varying standard deviations $\sqrt{h_{i,t}}$ on the i^{th} diagonal, i.e., $\mathbf{D}_t = diag\left\{\sqrt{h_{i,t}}\right\}$, and $\mathbf{Q}_t$ is the approximation of the conditional correlation matrix $\mathbf{R}_t$. In the DCC model, $\mathbf{Q}_t$ converges to the unconditional average correlation $\bar{\mathbf{R}} = \frac{1}{T}\sum \boldsymbol{\varepsilon}_{t-1}\boldsymbol{\varepsilon}_{t-1}'$, and $\Omega = (1-\alpha-\beta)\bar{\mathbf{R}}$. The positive semi-definiteness of $\mathbf{Q}_t$ is guaranteed if α and β are positive with $\alpha + \beta < 1$ and the initial matrix $\mathbf{Q}_1$ is positive definite.

Here, I estimate the conditional volatility $\mathbf{D}_t$ employing the univariate long memory volatility model of Zumbach (2006). I divide returns by their conditional volatility and use the standardized, zero-mean residuals $\boldsymbol{\varepsilon}_t = \mathbf{D}_t^{-1}\mathbf{r}_t$ to compute the quasi-conditional correlation matrix $\mathbf{Q}_t$. As the diagonal elements of $\mathbf{Q}_t$ are equal to unity only on average, $\mathbf{Q}_t$ is rescaled to obtain the conditional correlation matrix

$\mathbf{R}_t = diag\{\mathbf{Q}_t\}^{-1/2}\mathbf{Q}_t diag\{\mathbf{Q}_t\}^{-1/2}$. The conditional volatility $\mathbf{D}_t$ and conditional correlations $\mathbf{R}_t$ are then combined to estimate the conditional covariance matrix $\mathbf{H}_t$.

The *h*-step-ahead conditional covariance matrix is given by

$$\mathbf{H}_{t+h} = \mathbf{D}_{t+h}\mathbf{R}_{t+h}\mathbf{D}_{t+h}. \tag{5.13}$$

The forecast of each volatility in $\mathbf{D}_{t+h}$ is estimated using the forecast procedure derived by Zumbach (2006) (see Appendix 5.1 for the details). Since $\mathbf{R}_t$ is a non-linear process, the *h*-step forecast of $\mathbf{R}_t$ cannot be computed using a recursive procedure. However, assuming for simplicity that $E_t\left(\boldsymbol{\varepsilon}_{t+1}\boldsymbol{\varepsilon}'_{t+1}\right) \approx \mathbf{Q}_{t+1}$, Engle and Shephard (2001) show that the forecasts of $\mathbf{Q}_{t+h}$ and $\mathbf{R}_{t+h}$ are given by

$$\mathbf{Q}_{t+h} = \sum_{j=0}^{h-2}(1-\alpha-\beta)\overline{\mathbf{Q}}(\alpha+\beta)^j + (\alpha+\beta)^{h-1}\mathbf{Q}_{t+1}, \tag{5.14}$$

and

$$\mathbf{R}_{t+h} = diag\{\mathbf{Q}_{t+h}\}^{-\frac{1}{2}}\mathbf{Q}_{t+h} diag\{\mathbf{Q}_{t+h}\}^{-\frac{1}{2}}. \tag{5.15}$$

5.1.3 The FIGARCH(1,d,1)-DCC Model

In the FIGARCH(1,*d*,1) model of Baillie et al. (1996), the conditional volatility is modelled as:

$$h_t = \omega + [1-\beta L - (1-\phi L)(1-L)^d]\varepsilon_t^2 + \beta h_{t-1}. \tag{5.16}$$

Baillie et al. (1996) show that for $0 < d \le 1$, the FIGARCH process does not have finite unconditional variance, and is not weakly stationary, a feature shared with the IGARCH model. However, they show that the FIGARCH model is strictly stationary and ergodic by a direct extension of the corresponding proof for the IGARCH model.

The one-step ahead forecast of the FIGARCH(1,*d*,1) model is given by

$$h_{t+1} = \omega(1-\beta)^{-1} + [1-(1-\beta L)^{-1}(1-\phi L)(1-L)^d]\varepsilon_t^2, \tag{5.17}$$

and the *h*-step ahead forecast by

$$h_{t+h} = \omega(1-\beta)^{-1} + [1-(1-\beta L)^{-1}(1-\phi L)(1-L)^d]\varepsilon_{t+h-1}^2 . \quad (5.18)$$

To implement the FIGARCH(1,d,1) model in the multivariate context, I use the DCC approach described above, with the same forecast functions for $\mathbf{Q}_{t+h}$ and $\mathbf{R}_{t+h}$.

5.1.4 The CGARCH(1,1)-DCC Model

In the Component GARCH model of Engle and Lee (1999), the long memory volatility process h_t is modelled as the sum of a long term trend component, q_t, and a short term transitory component, s_t. The CGARCH(1,1) model has the following specification:

$$h_t - q_t = \alpha\left(\varepsilon_{t-1}^2 - q_{t-1}\right) + \beta(h_{t-1} - q_{t-1}) \quad (5.19)$$

$$q_t = \omega + \rho q_{t-1} + \phi(\varepsilon_{t-1}^2 - h_{t-1}) \quad (5.20)$$

where $s_t = h_t - q_t$ is the transitory volatility component. The volatility innovation $\varepsilon_{t-1}^2 - h_{t-1}$ drives both the trend and the transitory components. The long run component evolves over time following an AR process with ρ close to 1, while the short run component mean reverts to zero at a geometric rate $\alpha + \beta$. It is assumed that $0 < \alpha + \beta < \rho < 1$ so that the long run component is more persistent than the short run component.

The one-step ahead forecast of the CGARCH(1,1) model is given by

$$h_{t+1} = q_{t+1} + \alpha\left(\varepsilon_t^2 - q_t\right) + \beta(h_t - q_t) \quad (5.21)$$

$$q_{t+1} = \omega + \rho q_t + \phi(\varepsilon_t^2 - h_t), \quad (5.22)$$

and the h-step ahead forecast by

$$h_{t+h} = q_{t+h} + (\alpha+\beta)^{h-1}(h_t - q_t) \quad (5.23)$$

$$q_{t+h} = \frac{\omega}{1-\rho} + \rho^{h-1}\left(q_t - \frac{\omega}{1-\rho}\right). \quad (5.24)$$

As with the FIGARCH(1,d,1) model, in order to implement the CGARCH(1,1) model in the multivariate context, I use the DCC approach described above, with the same forecast functions for $\mathbf{Q}_{t+h}$ and $\mathbf{R}_{t+h}$.

5.1.5 The RiskMetrics EWMA Model

The short memory RiskMetrics EWMA covariance matrix is defined by

$$\mathbf{H}_t = \lambda \mathbf{H}_{t-1} + (1-\lambda)\boldsymbol{r}_{t-1}\boldsymbol{r}_{t-1}' \tag{5.25}$$

where λ is the decay factor $0 < \lambda < 1$. The larger the value of λ, the higher the persistence of the covariance matrix process and the lower the response of volatility to return shocks. It is straightforward to show that the h-step cumulative forecast of the EWMA model is given by

$$\mathbf{H}_{t+1:t+h} = h \times \mathbf{H}_{t+1}. \tag{5.26}$$

In the empirical analysis, I set λ to the values suggested by JP Morgan (1994) of 0.94 and 0.97 for daily and weekly forecasts, respectively.

5.1.6 The GARCH(1,1)-DCC Model

The short memory GARCH(1,1) model of Bollerslev (1990) is given by

$$h_t = \omega + \alpha\varepsilon_{t-1}^2 + \beta h_{t-1}. \tag{5.27}$$

The parameter α determines the speed at which the conditional variance responds to new information, while the parameter $\alpha + \beta$ determines how fast the conditional variance reverts to its long run average. In the GARCH(1,1) model, the weights on past squared errors decline at an exponential rate. The one-step ahead forecast of the GARCH(1,1) model is given by

$$h_{t+1} = \omega + \alpha\varepsilon_t^2 + \beta h_t, \tag{5.28}$$

and the h-step ahead forecast by

$$h_{t+h} = \sigma^2 + (\alpha+\beta)^{h-1}\left(h_{t+1} - \sigma^2\right) \tag{5.29}$$

where σ^2 is the unconditional variance. In order to implement the GARCH(1,1) model in the multivariate context, I again use the DCC approach described above, with the same forecast functions for $\mathbf{Q}_{t+h}$ and $\mathbf{R}_{t+h}$.

5.2 Forecast Performance Measures

The forecast performance of the six conditional volatility models is evaluated with a range of statistical and economic measures. I first measure the accuracy, bias and information content of the models' forecasts for each element of the covariance matrix using the squares and cross-products of daily returns as proxies for the actual variances and covariances being forecast. Forecast accuracy is evaluated using the Root Mean Squared Error (RMSE), the Mean Absolute Error (MAE) and the Heteroscedasticity-adjusted MSE (HMSE) of Bollerslev and Ghysels (1996). These are given by

$$RMSE = \sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(r_{i,t}r_{j,t} - \hat{\sigma}_{ij,t}\right)^2} \tag{5.30}$$

$$MAE = \frac{1}{T}\sum_{t=1}^{T}\left|r_{i,t}r_{j,t} - \hat{\sigma}_{ij,t}\right| \tag{5.31}$$

$$HMSE = \frac{1}{T}\sum_{t=1}^{T}\left(\frac{r_{i,t}r_{j,t}}{\hat{\sigma}_{ij,t}} - 1\right)^2 . \tag{5.32}$$

The Heteroscedasticity-adjusted MSE (HMSE) of Bollerslev and Ghysels (1996) penalises underpredictions more heavily than overpredictions, and hence may better match the user's actual loss function. Forecast bias and information content are measured using the Mincer-Zarnowitz regression, given by

$$r_{i,t}r_{j,t} = \alpha_{ij} + \beta_{ij}\hat{\sigma}_{ij,t} + \varepsilon_{ij} . \tag{5.33}$$

A forecast is conditionally unbiased (i.e. weak-form efficient) if and only if $\alpha_{ij} = 0$ and $\beta_{ij} = 1$.

As noted by Engle and Colacito (2006), the statistical evaluation of covariance matrix forecasts on an element-by-element basis has a number of drawbacks, particularly for high dimensional systems. In particular, direct comparisons between two covariance matrices are difficult because the distance between them is not well specified. Indeed, the statistical approaches described above implicitly assume that all elements of the covariance matrix are equally important (in the sense that the same error in each element is equally costly in economic terms), but there is no priori reason why this should necessarily be the case. Moreover, the use of low frequency realized volatility as a proxy for true volatility introduces considerable noise that inflates the forecast errors of the conditional volatility forecasts, substantially reducing their explanatory power. This has prompted tests of covariance matrix forecast performance based instead on economic loss criteria. Such tests have shown that conditional volatility models perform better when performance is measured using an economic loss function than when based on traditional statistical measures (see, for example, West et al., 1993, Engle et al., 1996).

In this chapter, I employ the economic loss function developed by Engle and Colacito (2006), who study the usefulness of forecasts of the conditional covariance matrix in an asset allocation framework. Assume that an investor allocates a fraction $\mathbf{w}_t$ of his wealth to *n* risky assets and the remainder $\left(1-\mathbf{w}_t^{'}\boldsymbol{\iota}\right)$ to the risk-free asset, where $\boldsymbol{\iota}$ is the $n\times 1$ unit vector. In the mean-variance optimization framework, the investor solves the following optimization problem at time *t*:

$$\min_{\mathbf{w}_t} \mathbf{w}_t^{'}\mathbf{H}_{t+1}\mathbf{w}_t \tag{5.34}$$

$$\text{subject to } \mathbf{w}_t^{'}\boldsymbol{\mu}+\left(1-\mathbf{w}_t^{'}\boldsymbol{\iota}\right)r_t^f = \mu_p^* \tag{5.35}$$

where $\mathbf{H}_{t+1}$ is the covariance matrix at time $t+1$, $\boldsymbol{\mu}$ is the vector of expected returns, r_t^f is the risk-free rate and μ_p^* is the target return. As $\boldsymbol{\mu}$ is assumed to be constant, the optimal weight of each asset changes over time as a result of changes in the covariance matrix. Since the true covariance matrix $\mathbf{H}_{t+1}$ is unobserved, the optimisation problem is solved using a forecast of $\mathbf{H}_{t+1}$ obtained from a multivariate conditional volatility model, to yield an approximation to the true optimal portfolio. The investor chooses among competing forecasts of the conditional covariance matrix on the basis of the volatility of the resulting portfolio. Engle and Colacito (2006) show that the lowest volatility of the investor's portfolio is obtained when the forecast covariance matrix is equal to the true covariance matrix, irrespective of both the expected excess return vector $\boldsymbol{\mu}$ and the target return μ_p^*. This then yields a straightforward economic test of the relative performance of competing covariance matrix forecasts based on the volatility of the optimal portfolio.

Engle and Colacito (2006) also note that in the bivariate context, the relative volatilities of portfolios depend on the relative returns of the *n* risky assets, and not on their absolute returns. Using polar coordinates, all possible pairs of relative expected returns can be expressed in the form $\mu = \left[\sin\frac{\pi j}{20}, \cos\frac{\pi j}{20}\right]$, for $j \in \{0, \ldots, 10\}$. When $j = 5$, for example, the expected returns are identical, which yields the global minimum variance portfolio. To obtain a single summary vector of expected returns, I construct prior probabilities for different vectors of expected returns using the sample data and the quasi-Bayesian approach introduced by Engle and Colacito (2006). I use these probabilities as weights to estimate a single weighted average vector of expected returns. Appendix 5.2 provides details on the derivation of these weights. In the empirical study, I assume a target excess return equal to 1.[9]

For each vector of expected returns, and for each pair of covariance matrix forecasts, I test whether the portfolio variances are equal using the Diebold and Mariano (1995) test. In particular, I consider the loss differential $u_t^k = \left(\sigma_t^{1,k}\right)^2 - \left(\sigma_t^{2,k}\right)^2$, where $\left(\sigma_t^{1,k}\right)^2$ and $\left(\sigma_t^{2,k}\right)^2$ are the conditional variances of portfolios 1 and 2, respectively, for the expected return vector μ^k. The null hypothesis is that the mean of u is equal to zero for all *k*. By regressing u_t^k on a constant, and using the Newey and West (1987) adjusted covariance matrix, the null hypothesis of equal variances is simply a test that the mean of u is zero. Engle and Colacito note that since u_t^k is itself heteroscedastic, a more efficient estimator can be obtained by dividing u by the true variance. Since the true covariance matrix is unknown and there are two estimators being compared, they suggest using the geometric mean of the two variance estimators as the denominator. The improved loss differential is given by

$$v_t^k = u_t^k \left[2\left(\mu^{k'}\left(\mathbf{H}_t^1\right)^{-1}\mu^k\right)\left(\mu^{k'}\left(\mathbf{H}_t^2\right)^{-1}\mu^k\right)\right]^{-1}. \tag{5.36}$$

I apply the Diebold and Mariano tests to both the *u* and *v* series. Joint tests for all vectors of expected returns are also conducted.

[9] The choice of the target return is immaterial in the sense that it does not affect the relative volatilities of portfolios.

5.3 Empirical Results

The empirical research employs the same three portfolios of assets as in Engle and Colacito ((2006) and one additional portfolio of the DJIA components. Details of the four datasets have been provided in Chapter 4. For each portfolio, the whole sample is divided into an initial estimation period of 252 observations (one year for the daily return series and five years for the weekly return series), and a forecast period of 5296, 895 and 4749 observations for the two bivariate portfolios, the international stock and bond portfolio and the DJIA component portfolio, respectively. The initial estimation period is used to estimate each model to generate out-of-sample forecasts of the covariance matrix for observation 253. The estimation window is then rolled forward one observation, the models re-estimated, and forecasts made for observation 254, and so on until the end of the sample is reached. I initially estimate the conditional covariance matrix using all multivariate conditional volatility models described in Section 5.1, except the FIGARCH(1,d,1)-DCC model. This model is excluded owing to the prohibitively short estimation period. In Section 5.3.4, I employ longer estimation periods and consider all six models.

5.3.1 Low Dimensional Systems: The Stock-Bond and S&P500-DJIA Portfolios

Statistical Evaluation

Table 5.1 reports the statistical evaluation of the accuracy of the five conditional volatility models using the RMSE, MAE, and HMSE measures for the two bivariate systems, namely the Stock-Bond and S&P500-DJIA portfolios. The LM-EWMA and LM-EWMA-DCC models yield identical RMSE, MAE and HMSE measures for the variances since in both models, the variance forecasts are based on the univariate long memory EWMA model. However, the LM-EWMA model performs better with respect to the covariance forecasts. The LM-EWMA model also yields the lowest RMSE and MAE for all elements in the Stock-Bond covariance matrix, while the short memory EWMA model performs best in the S&P500-DJIA case, although the difference between the EWMA and LM-EWMA models is small. Among the DCC models, the LM-EWMA-DCC model dominates, suggesting potential benefits from allowing for long memory in volatility. The short memory GARCH-DCC model is the worst model in terms of forecast accuracy under symmetric RMSE and MAE measures. However, the HMSE measure, which accounts for asymmetry in the treatment of under- and over-predictions, chooses the models least favoured by the RMSE and MAE measures, with the GARCH-DCC and CGARCH-DCC models producing the lowest forecast errors. I do not report the results of the HMSE for the low correlation Stock-Bond covariance because the conditional correlation for some individual observations is very close to zero, leading to very high values of $r_{i,t}r_{j,t}/\hat{\sigma}_{ij,t}$, which severely distorts the reported statistics.

The results of the Mincer-Zarnowitz regressions for the two bivariate systems are summarised in Table 5.2. The table reports the estimated coefficients of the regression, the R-squared statistic and the *p*-value for each element of the covariance matrix for the null hypothesis of conditional unbiasedness. The unbiasedness hypothesis cannot be rejected at conventional significance levels for any of the stock variance forecasts, nor for the covariance forecasts in the S&P500-DJIA system for the LM-EWMA and LM-EWMA-DCC models, but it is rejected in all other cases. In the cases that the unbiasedness hypothesis cannot be rejected, the LM-EWMA and LM-EWMA-DCC models have slope coefficients that are very close to unity. The EWMA model, though evidently not as efficient, performs slightly better in terms of explanatory power, as measured by the R-squared statistics. The CGARCH-DCC model performs rather badly, indeed only marginally better than the GARCH-DCC model.

Economic Evaluation

I use the forecasts of the covariance matrix to construct the minimum variance portfolios subject to a target excess return of 1. The relative conditional volatilities of portfolios constructed using the different conditional covariance matrix estimators and all possible vectors of expected returns are compared in Table 5.3. The pairs of Bayesian prior weighted returns are obtained from non-overlapping consecutive subsamples of 63 days (3 months) from the full datasets. Engle and Colacito (2006) show that by considering unconditional mean-adjusted returns, one can obtain a consistent estimator of the true conditional portfolio variance. The lowest conditional volatility, corresponding to the best covariance matrix estimate, is normalised to 100. The ‘Const’ portfolio is the fixed weight portfolio constructed with the ex-post unconditional covariance matrix. It is clear that the conditional covariance matrices generally outperform the unconditional covariance matrix, highlighting the economic value of volatility timing strategies. The results are favourable for the two LM-EWMA models. For both the low correlation Stock-Bond portfolio and the high correlation S&P500-DJIA portfolio, the LM-EWMA model consistently yields the lowest portfolio volatility. Incorporating long memory into the EWMA structure therefore appears to improve the forecasts of the conditional covariance matrix in a way that is economically valuable. Among the DCC models, the LM-EWMA-DCC model again dominates. Although the CGARCH model is designed to capture long memory volatility, its high degree of parameterisation, which potentially produces high estimation error, evidently hinders its performance. It is also interesting to note that the simple EWMA model outperforms more sophisticated models such as the GARCH-DCC and CGARCH-DCC models, and is even superior to the LM-EWMA-DCC models in most cases.

In practice, investors may be more concerned with out-of-sample realised volatility than conditional volatility. This is reported in Table 5.4 for each model for the two bivariate portfolios. Here, the results are similar, with the LM-EWMA model consistently yielding the lowest out-of-sample portfolio volatility.

Next, Diebold-Mariano tests are applied to test for the equality of different models with each vector of expected returns. Instead of reporting all of the results, I focus on those with expected returns close to the sample mean, i.e., $[\mu_{Stock}, \mu_{Bond}] = [0.95, 0.31]$ and $[0.99, 0.16]$, and $[\mu_{S\&P500}, \mu_{DJIA}] = [0.59, 0.81]$ and $[0.71, 0.71]$. Joint tests are also carried out for all vectors of expected returns applying the GMM method with a robust HAC covariance matrix. Table 5.5 and Table 5.6 show the results of both the standard and the improved Diebold-Mariano tests for the Stock-Bond and the S&P500-DJIA portfolios, respectively. Each cell in the tables corresponds to the test of the hypothesis that the two models in the row and column are equal in terms of volatility forecasting against the alternative that the model in the row is better or worse than the model in the column. A positive sign indicates that the model in the row is better than the model in the column, and vice-versa. The Diebold-Mariano tests confirm our earlier results. The conditional volatility models, especially the long memory volatility models, consistently outperform the unconditional constant model at conventional confidence levels in both versions of the Diebold-Mariano test. The standard Diebold-Mariano test also shows that the LM-EWMA model significantly dominates all other conditional volatility models, both short memory and long memory. With the improved version of the Diebold-Mariano test, the difference between each pair of models is less clearly marked and the outperformance of the LM-EWMA model is not significant in some cases. However, the Diebold-Mariano statistics are still positive.

5.3.2 High Dimensional Systems: The International Stock and Bond and the DJIA Portfolios

Economic Evaluation

In practice, a portfolio may comprise hundreds of assets and consequently an investor may want to examine the forecast performance of different conditional volatility models in a higher dimensional framework. In an asset allocation problem, the investor needs to estimate both the expected returns and the covariance matrix. However, since there are a prohibitively large number of possible expected return vectors for the high dimensional portfolios, I study the value of covariance matrix forecasts in two restricted cases. First, I form global minimum variance portfolios, where all expected returns are assumed to be equal. Note that the correctly specified covariance matrix will produce portfolios with the lowest volatility for any particular vector of expected returns, including the case that they are all equal. The results are reported in Table 5.7. For the multivariate portfolios, I assume a risk free rate of 4%. Consistent with the previous findings, in the international stock and bond portfolio, the LM-EWMA model yields the lowest conditional and out of sample volatilities. Owing to its simplicity, the simple EWMA model also performs very well, indeed better than the long memory LM-EWMA-DCC and CGARCH-DCC models. The short memory GARCH-DCC model is the least desirable model. However, the results for the DJIA portfolio are markedly different in that the DCC models tend to outperform the non-DCC models. Indeed, the superiority of the LM-EWMA model deteriorates significantly, although it still renders better forecasts than the EWMA model. Consistent with the results for the bivariate portfolios, the LM-EWMA-DCC model always produces the best portfolios among the DCC models.

In the second experiment, I form hedging portfolios in which one asset is hedged against all other assets in the portfolio. In so doing, I select the expected return vectors such that one entry is equal to one and all others are set to zero. With this strategy, the LM-EWMA-DCC model produces portfolios with the lowest conditional volatilities in 33 of the 34 hedging portfolios of international stocks and bonds, and 24 of the 29 portfolios of DJIA components (Table 5.8 and Table 5.9). The LM-EWMA model, though still dominating the EWMA model, is generally inferior to the GARCH-DCC and CGARCH-DCC models. The Diebold-Mariano joint tests for all hedging expected returns are applied and the findings are consistent with those of the relative volatilities (Table 5.10). The LM-EWMA-DCC model significantly outperforms all other models in both versions of the Diebold-Mariano tests. The LM-EWMA model performs badly, significantly outperforming only the EWMA model. In the DJIA portfolio, the LM-EWMA model is even dominated by the unconditional estimator. The result with the out-of-sample volatilities is similar and reported in Appendices 5.3 and 5.4.

These results show consistently that incorporating long memory in volatility dynamics improves the forecasts of the covariance matrix. Also, the more parsimonious the model, the less estimation error it generates and so the higher benefit it brings. The LM-EWMA model generally outperforms the EWMA model, while the LM-EWMA-DCC model always yields the best results among the DCC models. Besides, our results reveal an important difference in the relative forecasting power of the DCC and non-DCC models, in low dimensional and high dimensional systems. In particular, the greater flexibility that arises from separately estimating volatility and correlation is evidently beneficial in the high dimensional case. This deserves attention for future research.

5.3.3 Longer Horizon Forecasts

Practical problems often require forecasts over longer horizons than the one-step ahead forecasts considered above. In this section, I evaluate the forecast performance of different conditional volatility models, both statistically and economically, for horizons up to three months. Table 5.11 reports the RMSE of different conditional volatility models for one week, one month and one quarter ahead forecasts. The benchmarks are the true variances and covariances, proxied by the sum of squares and cross products of daily returns over the forecast horizons. The long memory volatility models generally outperform the short memory models, with the LM-EWMA models consistently yielding the smallest forecast errors, although the standard EWMA model again proves itself a simple yet statistically accurate model. Note again that the LM-EWMA and LM-EWMA-DCC models yield identical results for the variance forecasts since in both models, the variance forecasts are based on the univariate long memory EWMA model. The MAE results are similar and are reported in Appendix 5.5.

The Mincer-Zarnowitz regressions are implemented for the longer horizons in Table 5.12. Compared to the one-step ahead forecasts, the forecasts for longer horizons have higher information content, which may be attributable to the use of more accurate proxies of the true variances and covariances. Again, the two LM-EWMA models dominate the other short and long memory conditional volatility models at all forecast horizons. They are the only two models that generally yield conditionally unbiased forecasts for the elements of the covariance matrix.

The economic usefulness of alternative covariance matrix estimators is assessed for both low and high dimensional portfolios over longer investment horizons. I let the investor rebalance his portfolios weekly, monthly and quarterly. These rebalancing frequencies would cover the situations of most investors in practice, at least approximately, from a day trader to a mutual fund. Table 5.13 gives the out-of-sample performance of the weekly rebalanced bivariate portfolios. Results for the conditional volatilities are similar. The gains from employing the conditional volatility models of a weekly trader, as compared to those of a day trader, are smaller. The two LM-EWMA models still outperform both the short memory models and the long memory CGARCH-DCC model, though the gains, again, are lower. Among the two LM-EWMA models, neither dominates. The LM-EWMA model tends to perform better when the hypothetical vectors of expected returns are close to the unconditional mean and in the overall returns (which use the Bayesian priors as the weighting factors).

For the monthly and quarterly rebalanced portfolios, the results are similar. The two long memory EWMA models consistently produce better forecasts than the constant and short memory volatility models. The short memory conditional volatility models either rapidly revert to the unconditional volatility at an exponential rate or, in the case of the EWMA model, do not converge at all, and consequently have relatively uninteresting long-run forecasts. With slowly decaying autocorrelations, the long memory volatility models are able to better exploit past information and consequently yield more accurate forecasts over longer horizons. The outperformance of the two long memory models in the monthly and quarterly rebalanced portfolios confirms this intuition. To save space, only the out-of-sample results of the quarterly rebalanced portfolios are reported in Table 5.14. Monthly results are reported in Appendix 5.5.

Results for the two high dimensional portfolios are consistent with those for the two low dimensional portfolios. Under the global minimum variance strategy, the two long memory LM-EWMA models generally yield the most favourable results over horizons up to three months (Table 5.15). The CGARCH-DCC model also consistently outperforms the GARCH-DCC model. A similar conclusion follows from the hedging strategy. Table 5.16 reports the results of the Diebold-Mariano joint tests of the hedging DJIA portfolios for the equality of the different models' forecasts at different forecast horizons. Results, again, are in favour of the long memory volatility models. However, as with the daily rebalanced portfolios, the DCC models outperform the non-DCC models. The LM-EWMA model only dominates the EWMA model. For all rebalancing frequencies, the LM-EWMA-DCC model generally yields the most economically useful forecasts in both high dimensional portfolios. The Diebold-Mariano tests of the hedging international stock and bond portfolios are similar and reported in Appendix 5.6.

5.3.4 Additional Robustness Tests

Forecast performance is potentially affected by the size of the rolling window used to estimate the conditional volatility models. Therefore, I re-evaluate the forecast performance of the multivariate conditional volatility models using estimation windows of two years, five years and ten years of daily returns. In the cases of five-year and ten-year rolling windows, I also estimate the conditional covariance matrix using the FIGARCH-DCC model. I do not estimate the FIGARCH-DCC model with one-year and two-year rolling windows since the estimation of the FIGARCH model requires a prohibitively high upper lag cut-off. Following standard practice in the literature, I set the truncation lag for the FIGARCH model equal to 1000.

The outperformance of the two long memory volatility models reported above is found to be insensitive to the choice of estimation window length, in the both low dimension and high dimension cases. To save space, Table 5.17 only reports the economic evaluation for the two bivariate portfolios with a five-year estimation window. The two long memory models consistently produce forecasts that are more accurate and informative, and more economically useful than other short and long memory models. The simple EWMA model, although not as good as the LM-EWMA model, generally outperforms the more sophisticated GARCH model. The long memory FIGARCH model is the worst performing model, which may be attributable to the complexity of its specification. Although not reported, the use of long forecast horizons (one week, one month, and one quarter ahead) yields very similar conclusions. See Appendix 5.8 for more results.

Consistent results are identified with the multivariate portfolios. Table 5.18 reports the Diebold-Mariano tests for the hedging DJIA portfolios with five-year estimation window. The long memory models such as the LM-EWMA-DCC and CGARCH models consistently generate portfolios with the lowest variances in all investment horizons. Again, the DCC models outperform the non-DCC models in the multivariate systems. The high parameterisation of the FIGARCH model evidently hinders its performance, leading to poor results. More results of the two high dimensional portfolios with different estimation windows and rebalancing frequencies are reported in the Appendices 5.9 and 5.10.

5.4 Conclusion

In this chapter, I evaluate the economic benefits that arise from allowing for long memory in forecasting the covariance matrix of returns over both short and long horizons, using the asset allocation framework of Engle and Colacito (2006). In so doing, I compare the performance of a number of long memory and short memory multivariate volatility models. Incorporating long memory property improves forecasts of the conditional covariance matrix. In particular, I find that long memory volatility models generally dominate short memory and unconditional volatility models on the basis of both statistical and economic criteria, especially at longer horizons. Moreover, the relatively parsimonious long memory EWMA models outperform the more complex multivariate long memory GARCH models. The high degree of parameterisation of the Component GARCH and FIGARCH models evidently generates large estimation errors that are detrimental to their performance. The results are consistent across different datasets, and are robust to different investment horizons and estimation windows. The findings of the paper are consistent with those in the univariate volatility literature.

The non-DCC conditional covariance matrix estimators (such as the EWMA model with exponential weights and the LM-EWMA model with logarithmic weights) impose the same dynamic structure on all elements of the covariance matrix, which facilitates their implementation in high dimensional systems, but it comes at a cost in terms of estimation error. In a high dimensional system, employing a potentially less correctly specified but more flexible DCC structure may yield better results. Also, some of the eigenvalues of the high dimensional covariance matrix are inevitably very small, and so the inverse of the covariance matrix used in the asset allocation is likely to be ill-conditioned (see, for example, Zumbach, 2009a). This may partly explain the poor performance of the LM-EWMA model in large systems. It would be interesting to investigate this issue in greater detail.

The use of the long memory conditional covariance matrix produces optimal portfolios with lower realised volatility than the static unconditional covariance matrix. However, since our aim is simply to evaluate the forecasts of alternative conditional covariance matrices, and to choose the estimator that produces the lowest portfolio volatility, I do not explicitly consider realised portfolio returns. In particular, it does not follow that the portfolio with the lowest volatility is necessarily the best portfolio in terms of portfolio performance measures such as the Sharpe ratio. Thus it would also be of interest to investigate further the economic value of long memory volatility timing in the asset allocation framework, allowing for differences in return as well as risk, and for the effect of transaction costs.

Table 5.1. RMSE, MAE and HMSE for the Two Bivariate Systems

The table reports the RMSE, MAE and HMSE for each element of the conditional covariance matrix estimated using five multivariate conditional volatility models over the forecast period. The squares and cross-products of daily returns are used as proxies for the actual variances and covariances.

	EWMA	GARCH DCC	LM-EWMA	LM-EWMA DCC	CGARCH DCC
Panel A. Root Mean Square Error (RMSE)					
Variances					
Stock	4.7483	4.7953	4.7459	4.7459	4.7649
Bond	0.3964	0.3978	0.3957	0.3957	0.3988
S&P500	4.0336	4.0921	4.0434	4.0434	4.0646
DJIA	3.6876	3.7295	3.6900	3.6900	3.7076
Covariances					
Stock-Bond	0.7442	0.7593	0.7432	0.7536	0.7575
S&P500-DJIA	3.7951	3.8402	3.8015	3.8015	3.8166
Panel B. Mean Absolute Error (MAE)					
Variances					
Stock	1.5342	1.5372	1.5337	1.5337	1.5577
Bond	0.1803	0.1874	0.1799	0.1799	0.1880
S&P500	1.4088	1.4251	1.4089	1.4089	1.4407
DJIA	1.3079	1.3232	1.3077	1.3077	1.3357
Covariances					
Stock -Bond	0.3298	0.3335	0.3278	0.3295	0.3398
S&P500 -DJIA	1.3284	1.3435	1.3296	1.3306	1.3583
Panel C. Heteroskedasticity-adjusted Mean Square Error (HMSE)					
Variances					
Stock	13.5870	8.6617	11.3285	11.3285	8.9257
Bond	5.0709	4.1601	4.6114	4.6114	4.5520
S&P500	8.1334	5.5538	6.8703	6.8703	5.6967
DJIA	9.6358	5.9739	7.9015	7.9015	6.0413
Covariances					
S&P500-DJIA	9.7683	6.0861	8.0231	7.8463	6.1982

Table 5.2. Mincer–Zarnowitz Regressions for the Two Bivariate Systems

The table reports the estimated coefficients of the Mincer-Zarnowitz regression for each element of the covariance matrix. The *p*-values are for the tests of the joint hypothesis H_0: $\alpha_{ij} = 0$ and $\beta_{ij} = 1$. The numbers in the parentheses are the *t*-statistics to test $\alpha_{ij} = 0$ and $\beta_{ij} = 1$, respectively.

	Intercept	Slope	R^2	*p*-value	Intercept	Slope	R^2	*p*-value
	Panel A. EWMA				***Panel B. GARCH-DCC***			
Stock	0.167	0.888	0.176	0.421	0.312	0.783	0.169	0.003
	(1.218)	(-0.798)			(3.750)	(-2.993)		
Bond	0.058	0.657	0.037	0.000	0.051	0.642	0.030	0.000
	(5.443)	(-4.687)			(4.248)	(-4.656)		
S&P500	0.132	0.906	0.207	0.460	0.257	0.801	0.193	0.006
	(1.110)	(-0.739)			(3.423)	(-2.748)		
DJIA	0.142	0.889	0.179	0.381	0.242	0.801	0.168	0.009
	(1.275)	(-0.876)			(3.204)	(-2.627)		
Stock-Bond	-0.010	0.698	0.046	0.006	-0.052	0.554	0.022	0.001
	(-1.263)	(-2.762)			(-3.593)	(-3.472)		
S&P500-DJIA	0.130	0.899	0.195	0.430	0.228	0.811	0.183	0.012
	(1.180)	(-0.789)			(3.129)	(-2.490)		
	Panel C. LM-EWMA				***Panel D. LM-EWMA-DCC***			
Stock	-0.006	1.011	0.174	0.941	-0.006	1.011	0.174	0.941
	(-0.044)	(0.074)			(-0.044)	(0.074)		
Bond	0.050	0.706	0.037	0.000	0.050	0.706	0.037	0.000
	(4.377)	(-3.853)			(4.377)	(-3.853)		
S&P500	-0.005	1.011	0.201	0.933	-0.005	1.011	0.201	0.933
	(-0.044)	(0.084)			(-0.044)	(0.084)		
DJIA	0.005	1.001	0.175	0.991	0.005	1.001	0.175	0.991
	(0.040)	(0.011)			(0.040)	(0.011)		
Stock-Bond	-0.010	0.735	0.046	0.018	-0.051	0.648	0.030	0.002
	(-1.278)	(-2.357)			(-3.765)	(-3.105)		
S&P500-DJIA	-0.002	1.008	0.189	0.955	-0.010	1.012	0.189	0.927
	(-0.015)	(0.057)			(-0.087)	(0.092)		
	Panel E. CGARCH-DCC							
Stock	0.248	0.802	0.178	0.020				
	(2.712)	(-2.329)						
Bond	0.059	0.606	0.028	0.000				
	(4.916)	(-5.097)						
S&P500	0.208	0.817	0.203	0.030				
	(2.457)	(-2.172)						
DJIA	0.200	0.814	0.177	0.035				
	(2.411)	(-1.823)						
Stock-bond	-0.067	0.625	0.025	0.002				
	(-4.375)	(-2.952)						
S&P500-DJIA	0.184	0.822	0.192	0.042				
	(2.074)	(-2.035)						

Table 5.3. Comparison of Conditional Volatilities: Bivariate Portfolios

The table reports the average conditional volatilities for the two bivariate portfolios, constructed with the objective of minimizing variance subject to the target excess return of 1. Each row in the table shows the results for the pair of expected returns in the corresponding first two columns. The overall returns are the pair of weighted returns using the Bayesian prior probabilities. The lowest volatility in each row is normalised to 100.

Panel A. Stock-Bond Portfolio

μ_{Stock}	μ_{Bond}	**EWMA**	**GARCH DCC**	**LM EWMA**	**LM EWMA DCC**	**CGARCH DCC**	**Const**
0.00	1.00	100.247	102.156	100.000	101.909	104.206	105.938
0.16	0.99	100.243	102.330	100.000	101.983	104.557	105.148
0.31	0.95	100.402	102.580	100.000	102.077	105.025	104.422
0.45	0.89	100.505	102.587	100.000	101.956	105.457	103.754
0.59	0.81	100.580	102.465	100.000	101.827	105.655	103.074
0.71	0.71	100.521	102.632	100.000	101.746	105.472	102.840
0.81	0.59	100.390	102.705	100.000	102.017	104.974	103.507
0.89	0.45	100.317	102.673	100.000	102.040	104.832	105.564
0.95	0.31	100.237	101.994	100.000	101.301	104.207	108.904
0.99	0.16	100.465	101.949	100.000	100.509	103.752	111.038
1.00	0.00	100.385	103.277	100.000	102.335	105.200	106.434
Overall (weighted)		*100.208*	*102.097*	*100.000*	*101.496*	*104.307*	*108.365*

Panel B. S&P500-DJIA Portfolio

μ_{SP500}	μ_{DJIA}	**EWMA**	**GARCH DCC**	**LM EWMA**	**LM EWMA DCC**	**CGARCH DCC**	**Const**
0.00	1.00	100.402	102.054	100.089	100.000	102.233	101.206
0.16	0.99	100.491	101.925	100.076	100.000	102.227	101.435
0.31	0.95	100.511	101.534	100.090	100.000	102.136	101.865
0.45	0.89	100.308	100.836	100.022	100.000	101.695	102.751
0.59	0.81	100.172	100.917	100.000	100.559	101.547	105.244
0.71	0.71	100.191	102.031	100.000	100.658	102.946	102.088
0.81	0.59	100.184	101.040	100.000	100.658	101.317	107.019
0.89	0.45	100.377	101.256	100.000	100.586	101.988	106.425
0.95	0.31	100.347	101.473	100.000	100.404	102.137	103.928
0.99	0.16	100.256	101.681	100.000	100.219	102.119	102.558
1.00	0.00	100.261	101.697	100.000	100.087	102.045	101.784
Overall (weighted)		*100.204*	*102.057*	*100.000*	*100.719*	*103.022*	102.508

Table 5.4. Comparison of Out-of-Sample Volatilities: Bivariate Portfolios

The table reports the out-of-sample volatilities for the two bivariate portfolios, constructed with the objective of minimizing variance subject to the target excess return of 1. Each row in the table shows the results for the pair of expected returns in the corresponding first two columns. The overall returns are the pair of weighted returns using the Bayesian prior probabilities. The lowest volatility in each row is normalised to 100.

Panel A. Stock-Bond Portfolio

μ_{Stock}	μ_{Bond}	**EWMA**	**GARCH DCC**	**LM EWMA**	**LM EWMA DCC**	**CGARCH DCC**	**Const**
0.00	1.00	100.540	102.133	100.000	101.876	104.291	105.164
0.16	0.99	100.532	102.228	100.000	101.950	104.659	104.432
0.31	0.95	100.537	102.320	100.000	102.051	105.031	103.736
0.45	0.89	100.530	102.372	100.000	102.119	105.435	103.109
0.59	0.81	100.446	102.462	100.000	102.208	105.816	102.696
0.71	0.71	100.401	102.638	100.000	102.332	105.791	102.695
0.81	0.59	100.370	102.808	100.000	102.673	105.027	103.413
0.89	0.45	100.303	103.046	100.000	102.642	105.631	105.039
0.95	0.31	100.328	102.932	100.000	102.117	105.474	107.664
0.99	0.16	100.386	103.133	100.000	100.936	104.811	109.317
1.00	0.00	100.428	105.932	100.000	103.005	107.182	104.709
Overall (weighted)		*100.312*	*102.920*	*100.000*	*102.192*	*105.446*	*107.262*

Panel B. S&P500-DJIA Portfolio

μ_{SP500}	μ_{DJIA}	**EWMA**	**GARCH DCC**	**LM EWMA**	**LM EWMA DCC**	**CGARCH DCC**	**Const**
0.00	1.00	100.292	103.208	100.130	100.000	103.694	101.879
0.16	0.99	100.302	102.995	100.110	100.000	103.956	102.280
0.31	0.95	100.352	102.353	100.132	100.000	104.288	103.012
0.45	0.89	100.323	101.148	100.048	100.000	102.441	104.170
0.59	0.81	100.304	101.361	100.000	100.408	102.366	105.967
0.71	0.71	100.073	102.290	100.000	100.589	102.886	104.349
0.81	0.59	100.142	102.236	100.000	100.717	101.208	109.425
0.89	0.45	100.421	101.519	100.000	100.617	101.730	107.279
0.95	0.31	100.456	102.817	100.000	100.476	102.217	104.268
0.99	0.16	100.393	102.854	100.000	100.314	102.488	102.723
1.00	0.00	100.374	102.871	100.000	100.218	102.621	101.934
Overall (weighted)		*100.096*	*102.368*	*100.000*	*100.691*	*102.936*	*104.723*

Table 5.5. Diebold–Mariano Tests of the Stock-Bond Portfolio

The table reports the *t*-statistics of the Diebold-Mariano tests for the Stock-Bond portfolio using the improved version of the test described in Engle and Colacito (2006). Panels A and B correspond to [μ_{Stock}, μ_{Bond}] = [0.95; 0.31] and [0.99; 0.16], respectively. Panel C reports the joint tests of all the expected vectors of returns. The *t*-statistics for the standard version of the Diebold-Mariano test are reported in parentheses. A positive number indicates that the model in the row is better than that in the column, and vice-versa. *, ** and *** denote rejection of the equality hypothesis of the two models at 10%, 5% and 1% significance level.

	EWMA	GARCH DCC	LM EWMA	LM-EWMA DCC	CGARCH DCC	Const
Panel A. μ=[0.95; 0.31]						
EWMA		2.005*	-1.227	3.539***	2.763***	2.645***
		(2.746***)	(-2.151**)	(3.475***)	(4.757***)	(4.211***)
GARCH DCC	-2.005*		-2.302**	-0.952	1.132	1.632
	(-2.746***)		(-3.075**)	(-1.189)	(3.491***)	(2.469**)
LM-EWMA	1.227	2.302**		3.946***	3.175***	2.788***
	(2.151**)	(3.075***)		(4.042***)	(5.041***)	(4.391***)
LM-EWMA DCC	-3.539***	0.952	-3.946***		1.952**	2.102**
	(-3.475***)	(1.189)	(-4.042***)		(3.943***)	(3.094***)
CGARCH DCC	-2.763***	(-1.132)	-3.175***	-1.920*		1.065
	(-4.757***)	(-3.491***)	(-5.041***)	(-4.060***)		(1.382)
Constant	-2.645***	-1.632	-2.788***	-2.102**	-1.065	
	(-4.211***)	(-2.469**)	(-4.391***)	(-3.094***)	(-1.382)	
Panel B. μ=[0.99; 0.16]						
EWMA		0.804	-0.208	-0.287	1.023	2.272**
		(1.895*)	(-1.985**)	(-1.147)	(3.136***)	(4.266***)
GARCH DCC	-0.804		-0.812	-1.455	0.807	1.526
	(-1.895*)		(-2.087**)	(-1.723*)	(3.780***)	(2.673***)
LM-EWMA	0.208	0.812		-0.112	1.006	2.461**
	(1.985**)	(2.087**)		(1.766*)	(3.285***)	(4.499***)
LM-EWMA DCC	0.287	1.455	0.112		2.101**	2.341**
	(-1.147)	(1.723*)	(-1.766*)		(3.090***)	(3.886***)
CGARCH DCC	-1.023	-0.807	-1.006	-2.017**		1.560
	(-3.136***)	(-3.780***)	(-3.285***)	(-3.179***)		(1.954*)
Constant	-2.272**	-1.526	-2.461**	-2.366**	-1.560	
	(-4.266***)	(-2.673***)	(-4.499***)	(-3.945***)	(-1.954*)	
Panel C. Joint tests						
EWMA		2.001**	-0.422	1.741*	1.742*	3.229***
		(4.072***)	(-6.393***)	(5.806***)	(4.710***)	(7.849***)
GARCH DCC	-2.001**		-2.054**	-1.944*	0.171	0.044
	(-4.072***)		(-4.442***)	(-2.385**)	(4.719***)	(2.156**)
LM-EWMA	0.422	2.054**		1.807*	1.726*	3.629***
	(6.393***)	(4.442***)		(6.624***)	(4.957***)	(8.451***)
LM-EWMA DCC	-1.741*	1.944*	-1.807*		1.462	2.209**
	(-5.806***)	(2.385**)	(-6.624***)		(3.719***)	(5.003***)
CGARCH DCC	-1.742*	-0.171	-1.726*	-1.462		-0.322
	(-4.710***)	(-4.719***)	(-4.957***)	(-3.719***)		(-0.134)
Constant	-3.229***	-0.044	-3.629***	-2.209**	0.322	
	(-7.849***)	(-2.156**)	(-8.451***)	(-5.003***)	(0.134)	

Table 5.6. Diebold–Mariano Tests of the S&P500-DJIA Portfolio

The table reports the *t*-statistics of the Diebold-Mariano tests for the S&P500-DJIA portfolio using the improved version of the test described in Engle and Colacito (2006). Panel A and B correspond to [$\mu_{S\&P500}$, μ_{DJIA}] = [0.59; 0.81] and [0.71; 0.71], respectively. Panel C reports the joint tests of all the expected vectors of returns. The *t*-statistics for the standard version of the Diebold-Mariano test are reported in parentheses. A positive number indicates that the model in the row is better than that in the column, and vice-versa. *, ** and *** denote rejection of the equality hypothesis of the two models at 10%, 5% and 1% significance level.

	EWMA	GARCH DCC	LM EWMA	LM EWMA DCC	CGARCH DCC	Const
		Panel A. μ=[0.59; 0.81]				
EWMA		2.068**	-1.466	-0.073	1.617	3.452***
		(1.673*)	(-2.429**)	(0.156)	(2.157**)	(4.063***)
GARCH DCC	-2.068**		-2.434**	-2.204**	0.948	2.588***
	(-1.673*)		(-2.323**)	(-1.753*)	(1.048)	(3.613***)
LM EWMA	1.466	2.434**		0.352	1.687*	3.689***
	(2.429**)	(2.323**)		(0.911)	(2.401**)	(4.267***)
LM EWMA DCC	0.073	2.204**	-0.352		1.594	3.496***
	(-0.156)	(1.753*)	(-0.911)		(1.672*)	
CGARCH DCC	-1.617	-0.948	-1.687*	-1.594		1.335
	(-2.157**)	(-1.048)	(-2.401**)	(-1.672*)		(-2.413**)
Constant	-3.452***	-2.588***	-3.689***	-3.496***	-1.335	
	(-4.063***)	(-3.613***)	(-4.267***)	(-4.264***)	(-2.413**)	
		Panel B. μ=[0.71; 0.71]				
EWMA		2.373**	0.601	1.254	2.577***	2.036**
		(3.454***)	(-0.438)	(1.559)	(4.103***)	(2.497**)
GARCH DCC	-2.373**		-2.571***	-2.582***	-0.390	1.469
	(-3.454***)		(-4.046***)	(-3.449***)	(1.352)	(1.304)
LM EWMA	-0.601	2.571***		1.307	2.879***	2.161**
	(0.438)	(4.046***)		(2.145**)	(4.563***)	(2.615***)
LM EWMA DCC	-1.254	2.582***	-1.307		1.880*	2.152**
	(-1.559)	(3.449***)	(-2.145**)		(3.873***)	(2.481**)
CGARCH DCC	-2.577***	0.390	-2.879***	-1.880*		1.394
	(-4.103***)	(-1.352)	(-4.563***)	(-3.873***)		(0.844)
Constant	-2.036**	-1.469	-2.161**	-2.152**	-1.394	
	(-2.497**)	(-1.304)	(-2.615***)	(-2.481**)	(-0.844)	
		Panel C. Joint test				
EWMA		2.671***	0.338	1.246	3.179***	2.598***
		(5.166***)	(-3.329***)	(1.469)	(6.374***)	(5.222***)
GARCH DCC	-2.671***		-3.033***	-3.296***	-0.368	2.130**
	(-5.166***)		(-6.004***)	(-5.545***)	(1.307)	(3.795***)
LM EWMA	-0.338	3.033***		1.412	3.674***	2.769***
	(3.329***)	(6.004***)		(2.813***)	(7.161***)	(5.455***)
LM EWMA DCC	-1.246	3.296***	-1.412		2.611***	2.848***
	(-1.469)	(5.545***)	(-2.813***)		(5.939***)	(5.438***)
CGARCH DCC	-3.179***	0.368	-3.674***	-2.611***		2.004**
	(-6.374***)	(-1.307)	(-7.161***)	(-5.939***)		(3.161***)
Constant	-2.598***	-2.130**	-2.769***	-2.848***	-2.004**	
	(-5.222***)	(-3.795***)	(-5.455***)	(-5.438***)	(-3.161***)	

Table 5.7. Comparison of Volatilities: Multivariate Portfolios

The table reports the volatilities of the global minimum variance portfolios. The lowest volatility in each row is normalised to 100.

	EWMA	GARCH DCC	LM EWMA	LM EWMA DCC	CGARCH DCC	Const
Panel A. Conditional Volatilities						
International stock and bond portfolio	103.972	113.069	100.000	105.886	106.256	104.363
DJIA portfolio	126.605	102.881	106.554	100.000	102.478	132.890
Panel B. Out-of-Sample Volatilities						
International stock and bond portfolio	102.734	112.645	100.000	103.551	102.853	101.985
DJIA portfolio	125.268	104.119	105.393	100.000	103.663	135.269

Table 5.8. Comparison of Conditional Volatilities: Hedging International Portfolios

The table reports the average conditional volatilities of the minimum variance hedging portfolios subject to the target excess return of 1. Each row reports the results of the hedging portfolio, in which the asset of the country in the first column are hedged against all other assets. The lowest volatility is normalised to 100.

	EWMA	GARCH DCC	LM EWMA	LM-EWMA DCC	CGARCH DCC	Const
Panel A. International stocks						
Australia	112.154	103.993	108.319	100.000	103.661	114.440
Austria	113.783	104.576	109.120	100.000	104.900	125.200
Belgium	115.758	102.190	113.077	100.000	103.572	111.761
Canada	115.365	103.731	111.117	100.000	106.953	111.354
Denmark	115.634	101.428	111.750	100.000	101.463	101.470
France	118.771	100.915	115.481	100.000	101.461	109.469
Germany	115.396	105.132	110.583	100.000	105.142	114.046
Hongkong	116.388	102.980	110.743	100.000	103.102	109.661
Ireland	117.213	100.795	114.009	100.359	104.388	100.000
Italy	110.909	100.211	105.611	100.000	100.784	102.160
Japan	115.947	100.257	111.512	100.000	102.801	115.442
Mexico	114.938	102.112	111.307	100.000	103.330	106.627
Netherland	113.371	103.361	109.019	100.000	105.598	124.321
New Zealand	114.018	102.132	109.721	100.000	101.777	114.536
Norway	111.431	101.507	107.353	100.000	102.446	120.929
Singapore	111.810	102.789	107.644	100.000	101.854	105.794
Spain	118.453	103.257	114.973	100.000	101.266	119.418
Sweden	112.595	105.373	108.237	100.000	105.728	107.506
Switzerland	111.318	102.064	108.066	100.000	101.544	111.252
UK	111.234	104.059	107.062	100.000	101.788	114.067
US	114.647	107.630	109.543	100.000	109.473	111.657
Panel B. International bonds						
Austria	108.955	102.090	106.045	100.000	101.119	134.701
Belgium	117.297	105.712	112.631	100.000	104.344	137.973
Canada	115.487	102.466	111.641	100.000	102.217	106.356
Denmark	117.020	103.007	113.229	100.000	102.196	111.608
France	102.714	105.344	100.085	100.000	105.089	154.453
Germany	111.215	103.634	110.592	100.000	105.400	128.037
Ireland	111.222	100.582	106.525	100.000	103.367	135.702
Japan	114.813	101.211	110.368	100.000	100.572	103.768
Netherland	104.188	102.043	101.634	100.000	101.941	134.423
Sweden	109.796	101.580	106.949	100.000	101.816	103.997
Switzerland	106.535	101.430	104.057	100.000	100.249	111.290
UK	116.032	100.398	111.825	100.000	100.413	122.074
US	116.010	102.291	111.892	100.000	102.728	116.911

Table 5.9. Comparison of Conditional Volatilities: Hedging DJIA Portfolios

The table reports the average conditional volatilities of the minimum variance hedging portfolios subject to the target excess return of 1. Each row reports the results of the hedging portfolio, in which the asset of the country in the first column are hedged against all other assets. The lowest volatility is normalised to 100.

	EWMA	GARCH DCC	LM EWMA	LM-EWMA DCC	CGARCH DCC	Const
AA	138.806	100.865	113.661	100.000	101.612	105.519
AXP	137.429	101.766	112.050	100.000	101.318	107.559
BA	139.590	100.649	113.997	100.000	101.307	104.624
BAC	141.487	102.238	115.410	100.000	102.700	121.716
CAT	138.771	100.046	112.343	100.000	100.642	100.557
C	138.916	100.540	112.951	100.000	100.806	118.037
CVX	138.393	101.174	113.016	100.000	100.717	106.279
DD	136.763	100.476	112.285	100.000	100.421	112.715
DIS	137.711	100.000	111.968	100.169	100.482	101.953
GE	135.492	101.470	111.270	100.000	101.662	108.809
GM	138.716	101.470	112.822	100.000	101.663	113.654
HD	142.124	100.650	113.701	100.000	100.790	106.277
HPQ	141.407	101.845	114.473	101.516	102.421	100.000
IBM	141.815	100.188	114.008	100.000	100.570	101.147
INTC	137.506	100.366	112.643	100.000	100.664	100.114
JNJ	138.770	101.136	113.734	100.000	101.801	125.810
JPM	139.521	100.053	114.434	100.000	100.756	107.763
KO	139.793	100.359	114.269	100.000	100.850	104.120
MCD	140.553	100.287	113.798	100.287	101.194	100.000
MMM	141.341	103.563	115.321	102.832	104.116	100.000
MRK	140.490	101.020	114.149	100.000	101.563	110.373
MSFT	136.807	100.590	112.364	100.000	100.838	102.976
PFE	140.032	101.034	114.009	100.938	101.897	100.000
PG	138.747	100.455	114.054	100.000	100.762	101.742
T	140.041	100.496	114.605	100.000	100.903	103.782
UTX	134.319	100.318	110.637	100.000	101.551	104.866
VZ	137.122	100.977	112.456	100.000	100.771	108.897
WMT	140.294	100.505	114.417	100.000	101.028	106.758
XOM	135.238	100.462	111.311	100.000	100.409	111.073

Table 5.10. Diebold–Mariano Joint Tests: Hedging Multivariate Portfolios

The table reports the *t*-statistics of the Diebold–Mariano joint tests for the hedging multivariate portfolios, using the improved test of Engle and Colacito (2006). Panel A corresponds to the international stock and bond portfolio, while Panel B corresponds to the DJIA portfolio. The *t*-statistics for the standard test are reported in parentheses. A positive number indicates that the model in the row is better than the model in the column, and vice-versa. *, ** and *** denote rejection of the equality hypothesis of the two models at 10%, 5% and 1% significance level.

	EWMA	GARCH DCC	LM EWMA	LM-EWMA DCC	CGARCH DCC	Constant
Panel A. International Stock and Bond Portfolio						
EWMA		-3.44*** (-7.12***)	-3.62*** (-9.13***)	-4.63*** (-10.64***)	-3.49*** (-7.39***)	-3.04*** (-2.61***)
GARCH DCC	3.44*** (7.12***)		2.44*** (4.37***)	-2.78*** (-4.88***)	1.28 0.99	0.59 (5.06***)
LM-EWMA	3.62*** (9.13***)	-2.44*** (-4.37***)		-3.74*** (-7.51***)	-2.58*** (-4.53***)	-1.50 (-0.37)
LM-EWMA DCC	4.63*** (10.64***)	2.78*** (4.88***)	3.74*** (7.51***)	 (4.48***)	2.66*** (9.57***)	4.39***
CGARCH DCC	3.49*** (7.39***)	-1.28 (-0.99)	2.58*** (4.53***)	-2.66*** (-4.48***)		0.45 (3.97***)
Constant	3.04*** (2.61***)	-0.59 (-5.06***)	1.50 (0.37)	-4.39*** (-9.57***)	-0.45 (-3.97***)	
Panel B. DJIA Portfolio						
EWMA		-7.02*** (-37.75***)	-9.81*** (-40.43***)	-9.79*** (-36.79***)	-9.56*** (-38.16***)	-13.68*** (-32.09***)
GARCH DCC	7.02*** (37.75***)		1.77* (26.10***)	-1.52 (-4.10***)	-1.00 (1.28)	0.86 (11.10***)
LM-EWMA	9.81*** (40.43***)	-1.77* (-26.10***)		-8.79*** (-27.93***)	-7.58*** (-28.70***)	-6.57*** (-12.43***)
LM-EWMA DCC	9.79*** (36.79***)	1.52 (4.10***)	8.79*** (27.93***)		2.54** (6.32***)	7.18*** (13.77***)
CGARCH DCC	9.56*** (38.16***)	1.00 (-1.28)	7.58*** (28.70***)	-2.54** (-6.32***)		3.82*** (11.92***)
Constant	13.68*** (32.09***)	-0.86 (-11.10***)	6.57*** (12.43***)	-7.18*** (-13.77***)	-3.82*** (-11.92***)	

Table 5.11. RMSE for Longer Horizon Forecasts: Bivariate Systems

The table reports the RMSE for each element of the forecast conditional covariance matrix over the forecast period. The benchmarks are the realised variances and covariances, proxied by the sum of squares and cross products of returns over the forecast horizons, respectively.

	EWMA	GARCH DCC	LM EWMA	LM-EWMA DCC	CGARCH DCC
Panel A. One Week (5-Step) ahead Forecasts					
Variances					
Stock	12.675	13.207	12.668	12.668	12.330
Bond	0.918	0.938	0.901	0.901	0.952
S&P500	10.308	10.676	10.411	10.411	10.039
DJIA	9.611	9.892	9.638	9.638	9.286
Covariances					
Stock-Bond	1.814	1.892	1.788	1.811	1.880
S&P500-DJIA	9.783	10.062	9.851	9.858	9.460
Panel B. One Month (21-Step) ahead Forecasts					
Variances					
Stock	49.667	54.759	47.789	47.789	51.491
Bond	2.230	2.348	2.128	2.128	2.487
S&P500	42.384	47.684	41.015	41.015	44.696
DJIA	38.029	40.468	36.789	36.789	38.583
Covariances					
Stock-Bond	4.523	4.875	4.737	4.691	4.536
S&P500-DJIA	39.655	43.093	38.330	38.350	40.755
Panel C. One Quarter (63-Step) ahead Forecasts					
Variances					
Stock	151.499	168.737	146.514	146.514	165.768
Bond	5.348	5.630	4.983	4.983	5.974
S&P500	133.140	151.082	129.748	129.748	142.108
DJIA	113.798	128.904	111.416	111.416	120.125
Covariances					
Stock-Bond	7.729	10.210	7.893	10.470	9.377
S&P500-DJIA	121.775	137.591	118.879	118.890	128.746

Table 5.12. Mincer–Zarnowitz Regressions for Longer Horizons: Bivariate Systems

The table reports the results of the Mincer-Zarnowitz regressions for longer horizon forecasts of each element of the covariance matrix. The *p*-values are for the tests of the joint hypothesis: H_0: $\alpha_{ij} = 0$ and $\beta_{ij} = 1$.

	EWMA				LM-EWMA			
	Intercept	**Slope**	**R^2**	***p*-value**	**Intercept**	**Slope**	**R^2**	***p*-value**
Panel A. One Week (5-Step) ahead Forecasts								
Stock	1.042	0.860	0.412	0.000	0.074	0.997	0.408	0.025
Bond	0.282	0.680	0.161	0.000	0.199	0.782	0.170	0.000
S&P500	0.759	0.895	0.497	0.002	-0.051	1.019	0.480	0.341
DJIA	0.839	0.870	0.437	0.000	0.049	0.999	0.426	0.021
Stock-Bond	-0.066	0.686	0.176	0.000	-0.061	0.752	0.183	0.000
S&P500-DJIA	0.758	0.884	0.472	0.000	-0.011	1.011	0.456	0.133
Panel B. One Month (21-Step) ahead Forecasts								
Stock	8.454	0.730	0.322	0.000	5.739	0.825	0.348	0.103
Bond	0.931	0.764	0.350	0.003	0.615	0.871	0.389	0.148
S&P500	6.408	0.784	0.394	0.005	3.996	0.875	0.413	0.325
DJIA	6.708	0.748	0.346	0.001	4.459	0.836	0.365	0.129
Stock-Bond	-0.297	0.776	0.379	0.034	-0.344	0.739	0.326	0.001
S&P500-DJIA	6.259	0.768	0.371	0.003	3.968	0.860	0.391	0.236
Panel C. One Quarter (63-Step) ahead Forecasts								
Stock	44.494	0.536	0.172	0.005	37.168	0.623	0.167	0.104
Bond	3.462	0.660	0.468	0.000	2.641	0.754	0.482	0.041
S&P500	34.961	0.614	0.244	0.014	28.389	0.697	0.226	0.192
DJIA	34.137	0.583	0.221	0.008	29.114	0.647	0.196	0.094
Stock-Bond	-0.714	0.886	0.669	0.422	-0.756	0.919	0.649	0.840
S&P500-DJIA	33.000	0.600	0.232	0.011	27.280	0.677	0.221	0.150

	GARCH-DCC				LM-EWMA-DCC				CGARCH-DCC			
	Intercept	**Slope**	**R^2**	***p*-value**	**Intercept**	**Slope**	**R^2**	***p*-value**	**Intercept**	**Slope**	**R^2**	***p*-value**
Panel A. One Week (5-Step) ahead Forecasts												
Stock	1.751	0.750	0.401	0.000	0.074	0.997	0.408	0.025	1.139	0.805	0.467	0.000
Bond	0.251	0.664	0.120	0.000	0.199	0.782	0.170	0.000	0.305	0.602	0.111	0.000
S&P500	1.239	0.806	0.480	0.000	-0.051	1.019	0.480	0.341	0.704	0.870	0.528	0.000
DJIA	1.237	0.796	0.423	0.000	0.049	0.999	0.426	0.021	0.631	0.872	0.477	0.000
Stock-Bond	-0.184	0.610	0.122	0.000	-0.153	0.730	0.170	0.000	-0.230	0.654	0.127	0.000
S&P500-DJIA	1.134	0.811	0.457	0.000	-0.041	1.015	0.456	0.200	0.586	0.878	0.508	0.000
Panel B. One Month (21-Step) ahead Forecasts												
Stock	13.066	0.564	0.305	0.000	5.739	0.825	0.348	0.103	10.345	0.615	0.371	0.000
Bond	0.819	0.747	0.274	0.008	0.615	0.871	0.389	0.148	1.254	0.624	0.236	0.000
S&P500	11.133	0.604	0.342	0.000	3.996	0.875	0.413	0.325	8.775	0.652	0.411	0.000
DJIA	9.431	0.629	0.330	0.000	4.459	0.836	0.365	0.129	7.609	0.667	0.383	0.000
Stock-Bond	-0.543	0.662	0.335	0.000	-0.459	0.702	0.367	0.000	-0.623	0.770	0.389	0.000
S&P500-DJIA	9.708	0.625	0.339	0.000	3.935	0.862	0.390	0.255	7.670	0.667	0.401	0.000
Panel C. One Quarter (63-Step) ahead Forecasts												
Stock	60.382	0.336	0.065	0.000	37.168	0.623	0.167	0.104	54.682	0.381	0.089	0.000
Bond	2.910	0.687	0.353	0.016	2.641	0.754	0.482	0.041	2.954	0.652	0.294	0.004
S&P500	50.654	0.414	0.098	0.000	28.389	0.697	0.226	0.192	39.468	0.525	0.152	0.007
DJIA	49.335	0.379	0.086	0.000	29.114	0.647	0.196	0.094	38.228	0.505	0.143	0.004
Stock-Bond	-1.326	0.721	0.479	0.006	-0.975	0.672	0.487	0.000	-1.495	0.794	0.542	0.036
S&P500-DJIA	47.863	0.401	0.091	0.000	27.294	0.678	0.210	0.155	36.757	0.524	0.148	0.008

Table 5.13. Comparison of Out-of-Sample Volatilities: Bivariate Portfolios with Weekly Rebalancing Frequency

The table reports the out-of-sample volatilities for the weekly rebalanced bivariate portfolios, constructed with the objective of minimizing variance subject to the target excess return of 1. Each row in the table reports the results for the pair of expected returns in the corresponding first two columns. The overall returns are the pair of weighted returns using the Bayesian prior probabilities. The lowest volatility in each row is normalised to 100.

Panel A. Stock-Bond Portfolio

μ_{Stock}	μ_{Bond}	EWMA	GARCH DCC	LM EWMA	LM EWMA DCC	CGARCH DCC	Const
0.00	1.00	100.472	100.590	100.000	100.000	102.243	105.903
0.16	0.99	100.524	100.524	100.058	100.000	102.446	105.533
0.31	0.95	100.562	100.449	100.169	100.000	102.528	105.056
0.45	0.89	100.635	100.424	100.265	100.000	102.594	104.553
0.59	0.81	100.730	100.389	100.292	100.000	102.530	103.990
0.71	0.71	100.652	100.478	100.217	100.000	102.262	103.306
0.81	0.59	100.528	100.642	100.000	100.000	101.925	102.491
0.89	0.45	100.447	100.767	100.000	100.032	101.374	102.204
0.95	0.31	100.376	101.207	100.000	100.107	102.468	103.165
0.99	0.16	100.670	101.547	100.000	100.485	103.486	105.149
1.00	0.00	100.921	101.800	100.000	100.879	102.550	105.271
Overall (weighted)		*100.360*	*101.099*	*100.000*	*100.049*	*102.178*	*102.913*

Panel B. S&P500-DJIA Portfolio

μ_{SP500}	μ_{DJIA}	EWMA	GARCH DCC	LM EWMA	LM EWMA DCC	CGARCH DCC	Const
0.00	1.00	100.621	100.966	100.207	100.000	100.897	101.655
0.16	0.99	100.762	100.645	100.352	100.000	100.821	102.345
0.31	0.95	100.941	100.141	100.612	100.000	100.612	103.388
0.45	0.89	101.749	100.000	101.435	100.700	100.770	105.493
0.59	0.81	100.842	100.000	100.591	100.364	100.682	105.755
0.71	0.71	100.184	100.811	100.000	100.199	100.719	101.560
0.81	0.59	100.431	101.940	100.000	100.367	101.574	107.827
0.89	0.45	100.396	100.759	100.000	100.231	100.660	105.805
0.95	0.31	100.400	101.422	100.044	100.000	101.244	102.222
0.99	0.16	100.278	101.893	100.111	100.000	101.336	100.557
1.00	0.00	100.331	102.118	100.199	100.132	101.390	100.000
Overall (weighted)		*100.225*	*100.856*	*100.000*	*100.187*	*100.677*	*101.613*

Table 5.14. Comparison of Out-of-Sample Volatilities: Bivariate Portfolios with Quarterly Rebalancing Frequency

The table reports the out-of-sample volatilities for the quarterly rebalanced bivariate portfolios, constructed with the objective of minimizing variance subject to the target excess return of 1. Each row in the table reports the results for the pair of expected returns of the corresponding first two columns. The overall returns are the pair of weighted returns using the Bayesian prior probabilities. The lowest volatility in each row is normalised to 100.

Panel A. Stock-Bond Portfolio

μ_{Stock}	μ_{Bond}	EWMA	GARCH DCC	LM EWMA	LM EWMA DCC	CGARCH DCC	Const
0.00	1.00	100.622	100.000	100.000	100.000	100.415	110.581
0.16	0.99	101.031	100.000	100.412	100.000	100.412	109.691
0.31	0.95	101.603	100.000	100.601	100.000	100.200	108.417
0.45	0.89	101.908	100.000	100.954	100.191	100.000	107.252
0.59	0.81	102.135	100.000	101.246	100.534	100.000	106.228
0.71	0.71	101.789	100.650	101.138	100.813	100.000	105.528
0.81	0.59	100.728	101.019	100.146	100.146	100.000	104.803
0.89	0.45	100.783	102.611	100.000	100.261	101.044	106.658
0.95	0.31	100.818	103.855	100.467	100.000	102.336	109.813
0.99	0.16	100.526	103.891	100.315	100.000	103.260	114.826
1.00	0.00	100.000	102.844	100.267	101.422	101.778	113.867
Overall (weighted)		*100.727*	*103.749*	*100.336*	*100.000*	*102.098*	*109.219*

Panel B. S&P500-DJIA Portfolio

μ_{SP500}	μ_{DJIA}	EWMA	GARCH DCC	LM EWMA	LM EWMA DCC	CGARCH DCC	Const
0.00	1.00	100.000	102.597	100.519	100.519	103.377	101.818
0.16	0.99	100.000	102.649	100.442	100.442	103.311	101.987
0.31	0.95	100.000	102.482	100.355	100.177	102.837	102.128
0.45	0.89	100.000	101.713	100.000	100.000	101.976	102.635
0.59	0.81	100.629	101.258	100.449	100.000	101.797	105.121
0.71	0.71	101.020	100.960	100.000	101.621	106.363	102.461
0.81	0.59	100.000	106.854	101.168	102.103	106.464	109.891
0.89	0.45	100.000	105.181	100.361	101.084	105.301	108.434
0.95	0.31	100.163	103.431	100.000	100.490	104.739	104.248
0.99	0.16	100.206	102.263	100.000	100.206	103.909	102.469
1.00	0.00	100.000	101.716	100.000	100.000	103.676	101.471
Overall (weighted)		*100.707*	*100.865*	*100.000*	*101.623*	*106.049*	*102.937*

Table 5.15. Comparison of Volatilities: Multivariate Portfolios with Different Rebalancing Frequencies

The table reports the out-of-sample volatilities of the global minimum variance multivariate portfolios. Conditional volatilities are reported in parentheses. The lowest volatility in each row is normalised to 100.

	EWMA	GARCH DCC	LM EWMA	LM-EWMA DCC	CGARCH DCC
Panel A. International Stock and Bond Portfolio					
Monthly rebalancing	102.884 (102.502)	123.706 (113.730)	100.000 (100.000)	104.699 (105.481)	110.065 (111.344)
Quarterly rebalancing	106.339 (103.374)	128.935 (116.535)	100.000 (100.000)	105.900 (107.679)	113.939 (110.481)
Panel B. DJIA Portfolio					
Weekly rebalancing	107.748 (106.180)	102.985 (100.444)	104.699 (102.825)	100.000 (100.000)	102.053 (100.597)
Monthly rebalancing	105.261 (107.272)	103.892 (104.615)	103.612 (104.311)	100.000 (100.000)	102.227 (104.273)
Quarterly rebalancing	115.392 (114.665)	107.737 (108.247)	108.621 (107.377)	100.000 (100.000)	104.939 (103.760)

Table 5.16. Diebold–Mariano Joint Tests: Hedging DJIA Portfolios with Different Rebalancing Frequencies

The table reports the *t*-statistics of the Diebold–Mariano joint tests for the hedging DJIA portfolios, using the improved test of Engle and Colacito (2006). The *t*-statistics for the standard test are reported in parentheses. A positive number indicates that the model in the row is better than the model in the column, and vice-versa. *, ** and *** denote rejection of the equality hypothesis of the two models at 10%, 5% and 1% significance level.

	EWMA	GARCH DCC	LM-EWMA	LM-EWMA DCC	CGARCH DCC
Panel A. Weekly Rebalancing					
EWMA		-5.02***	-5.94***	-6.81***	-6.24***
		(-15.88***)	(-13.91***)	(-16.71***)	(-17.80***)
GARCH-DCC	5.02***		4.15***	-2.17**	0.07
	(15.88***)		(13.92***)	(-2.29**)	(2.80***)
LM-EWMA	5.94***	-4.15***		-6.28***	-5.24***
	(13.91***)	(-13.92***)		(-15.15***)	(-15.78***)
LM-EWMA-DCC	6.81***	2.17**	6.28***		3.20***
	(16.71***)	(2.29**)	(15.15***)		(5.10***)
CGARCH-DCC	6.24***	-0.07	5.24***	-3.20***	
	(17.80***)	(-2.80***)	(15.78***)	(-5.10***)	
Panel B. Monthly Rebalancing					
EWMA		-4.25***	-6.76***	-5.34***	-4.63***
		(-11.32***)	(-14.94***)	(-11.69***)	(-10.75***)
GARCH-DCC	4.25***		3.25***	0.13	1.58
	(11.32***)		(7.73***)	(-0.64)	(1.28)
LM-EWMA	6.76***	-3.25***		-4.28***	-3.29***
	(14.94***)	(-7.73***)		(-8.40***)	(-7.18***)
LM-EWMA-DCC	5.34***	-0.13	4.28***		1.72*
	(11.69***)	(0.64)	(8.40***)		(2.27**)
CGARCH-DCC	4.63***	-1.58	3.29***	-1.72*	
	(10.75***)	(-1.28)	(7.18***)	(-2.27**)	
Panel C. Quarterly Rebalancing					
EWMA		-1.33	-6.96***	-6.07***	-4.97***
		(-10.27***)	(-12.64***)	(-10.78***)	(-10.29***)
GARCH-DCC	1.33		-0.04	-0.92	-0.70
	(10.27***)		(6.25***)	(-1.16)	(0.52)
LM-EWMA	6.96***	0.04		-4.34***	-2.75***
	(12.64***)	(-6.25***)		(-7.10***)	(-6.11***)
LM-EWMA-DCC	6.07***	0.92	4.34***		0.17
	(10.78***)	(1.16)	(7.10***)		1.59
CGARCH-DCC	4.97***	0.70	2.75***	-0.17	
	(10.29***)	(-0.52)	(6.11***)	-1.59	

Table 5.17. Comparison of Out-of-Sample Volatilities: Bivariate Portfolios with 5-Year Estimation Window

The table reports the out-of-sample volatilities for the minimum variance bivariate portfolios, constructed using 5-year estimation window and subject to the excess target return of 1. Each row in the table reports the results for the pair of expected returns in the corresponding first two columns. The overall returns are the pair of weighted returns using the Bayesian prior probabilities. The lowest volatility in each row is normalised to 100.

Panel A. Stock-Bond Portfolio

μ_{Stock}	μ_{Bond}	EWMA	GARCH DCC	LM EWMA	LM EWMA DCC	CGARCH DCC	FIGARCH DCC	Const
0.00	1.00	100.564	102.177	100.000	102.075	103.637	106.967	105.610
0.16	0.99	100.608	102.128	100.000	102.103	103.572	108.715	104.890
0.31	0.95	100.588	102.009	100.000	102.058	103.552	110.828	104.140
0.45	0.89	100.556	101.922	100.000	101.992	103.659	112.969	103.474
0.59	0.81	100.471	101.883	100.000	101.904	103.808	114.955	102.974
0.71	0.71	100.387	101.971	100.000	101.817	103.923	116.293	103.015
0.81	0.59	100.290	102.170	100.000	101.811	103.861	116.761	104.032
0.89	0.45	100.162	102.184	100.000	101.697	103.424	116.293	106.272
0.95	0.31	100.236	101.972	100.000	101.191	103.523	114.823	109.278
0.99	0.16	100.395	101.376	100.000	100.617	103.177	111.655	110.229
1.00	0.00	100.427	102.484	100.000	102.902	104.224	104.115	104.567
Overall (weighted)		*100.207*	*102.010*	*100.000*	*101.296*	*103.446*	*115.106*	*108.884*

Panel B. S&P500-DJIA Portfolio

μ_{SP500}	μ_{DJIA}	EWMA	GARCH DCC	LM EWMA	LM EWMA DCC	CGARCH DCC	FIGARCH DCC	Const
0.00	1.00	100.217	101.485	100.031	100.000	101.671	158.014	101.702
0.16	0.99	100.210	101.446	100.000	100.000	101.683	165.843	102.182
0.31	0.95	100.254	101.458	100.021	100.000	101.733	170.716	103.042
0.45	0.89	100.282	101.408	100.000	100.047	101.847	145.211	104.570
0.59	0.81	100.226	101.068	100.000	100.411	100.503	139.945	107.094
0.71	0.71	100.102	101.108	100.000	100.382	101.439	111.423	104.680
0.81	0.59	100.098	101.238	100.000	100.543	100.285	117.528	110.305
0.89	0.45	100.469	101.278	100.000	100.540	101.293	118.835	107.841
0.95	0.31	100.528	101.468	100.000	100.372	101.566	129.928	104.463
0.99	0.16	100.470	101.510	100.000	100.198	101.659	136.940	102.748
1.00	0.00	100.413	101.535	100.000	100.118	101.712	139.327	101.860
Overall (weighted)		*100.123*	*101.209*	*100.000*	*100.448*	*101.403*	*112.275*	*105.138*

Table 5.18. Diebold–Mariano Joint Tests: Hedging DJIA Portfolios with 5-Year Estimation Window

The table reports the *t*-statistics of the Diebold–Mariano tests for the hedging DJIA portfolios, using the improved test of Engle and Colacito (2006). The *t*-statistics for the standard test are reported in parentheses. A positive number indicates that the model in the row is better than the model in the column, and vice-versa. *, ** and *** denote rejection of the equality hypothesis of the two models at 10%, 5% and 1% significance level.

	EWMA	GARCH DCC	LM-EWMA	LM-EWMA DCC	CGARCH DCC	FIGARCH DCC	EWMA	GARCH DCC	LM-EWMA	LM-EWMA DCC	CGARCH DCC	FIGARCH DCC
Panel A. Daily Rebalancing							***Panel B. Weekly Rebalancing***					
EWMA		-8.03***	-9.50***	-8.43***	8.73***	1.60		-6.61***	-6.98***	-6.24***	-7.14***	2.67***
		(-34.3***)	(-36.9***)	(-33.6***)	(34.4***)	(-13.3***)		(-19.9***)	(-18.3***)	(-19.3***)	(-19.8***)	(-3.91***)
GARCH-DCC	9.50***		7.33***	-3.18***	-0.89	3.4***	6.61***		5.58***	-2.33**	-1.21	3.74***
	(36.9***)		(26.6***)	(-6.09***)	(-0.74)	(7.79***)	(19.9***)		(16.35***)	(-2.91***)	(-0.62)	(7.66***)
LM-EWMA	8.43***	-7.33***		-7.81***	-8.96***	2.79***	6.98***	-5.58***		-5.47***	-6.02***	2.96***
	(33.6***)	(-26.6***)		(-25.8***)	(-27.4***)	(1.44)	(18.3***)	(-16.4***)		(-15.8***)	(-16.2***)	(-0.34)
LM-EWMA-DCC	8.73***	3.18***	7.81***		2.12**	3.49***	6.24***	2.33**	5.47***		1.72*	3.92***
	(34.4***)	(6.09***)	(25.8***)		(4.55***)	(7.93***)	(19.3***)	(2.91***)	(15.78***)		(2.87***)	(7.64***)
CGARCH-DCC	8.03***	0.89	8.96***	-2.12**		3.56***	7.14***	1.21	6.02***	-1.72*		3.85***
	(34.3***)	(0.74)	(27.44***)	(-4.55***)		(7.79***)	(19.8***)	(0.62)	(16.17***)	(-2.87***)		(7.45***)
FIGARCH-DCC	-1.60	-3.4***	-2.79***	-3.49***	-3.56***		-2.67***	-3.74***	-2.96***	-3.92***	-3.85***	
	(13.3***)	(-7.79***)	(-1.44)	(-7.93***)	(-7.79***)		(3.91***)	(-7.66***)	(0.34)	(-7.64***)	(-7.45***)	
Panel C. Monthly Rebalancing							***Panel D. Quarterly Rebalancing***					
EWMA		-4.12***	-7.23***	-5.17***	-7.04***	0.29		-3.12***	-5.85***	-4.72***	-4.11***	1.46
		(-11.6***)	(-15.9***)	(-12.6***)	(-13.2***)	(-5.25***)		(-8.67***)	(-10.7***)	(-9.08***)	(-8.76***)	(1.17)
GARCH-DCC	4.12***		3.10***	-0.44	0.76	2.55***	3.12***		1.77	0.48	0.67	1.46
	(11.6***)		(-.69***)	(-0.87)	(0.86)	(4.54***)	(8.67***)		(5.41***)	(0.59)	(1.29)	(1.77*)
LM-EWMA	7.23***	-3.10***		-4.22***	-5.83***	1.39	5.85***	-1.77		-3.19***	-2.62***	1.48
	(15.9***)	(-7.69***)		(-8.84***)	(-9.23***)	(-1.49)	(10.7***)	(-5.41***)		(-5.93***)	(-5.48***)	(1.53)
LM-EWMA-DCC	5.17***	0.44	4.22***		1.33	2.68***	4.72***	-0.48	3.19***		0.40	1.51
	(12.6***)	(0.87)	(8.84***)		(1.95*)	(4.25***)	(9.08***)	(-0.59)	(5.93***)		(0.94)	(1.76*)
CGARCH-DCC	7.04***	-0.76	5.83***	-1.33		2.30***	4.11***	-0.67	2.62***	-0.40		1.50
	(13.2***)	(-0.86)	(9.23***)	(-1.95*)		(3.54***)	(8.76***)	(-1.29)	(5.48***)	(-0.94)		(1.75*)
FIGARCH-DCC	-0.29	-2.55***	-1.39	-2.68***	-2.30***		-1.46	-1.46	-1.48	-1.51	-1.50	
	(5.25***)	(-4.54***)	(1.49)	(-4.25***)	(-3.54***)		(-1.17)	(-1.77*)	(-1.53)	(-1.76*)	(-1.75*)	

Appendix 5.1. LM-EWMA Conditional Covariance Matrix Forecasts

Volatility Forecast

Zumbach (2006) derives the recursive formula to forecast long memory volatility. Employing the processes in (5.3) and (5.4), the *j*-step-ahead conditional volatility, given information at time *t*, can be derived from the system of recursive conditional equations:

$$E_t(\sigma^2_{t+j}) = \sum_{k=1}^{K} w_k E_t(\sigma^2_{t+j-1,k}) \tag{1}$$

$$E_t\left(\sigma^2_{t+j-1,k}\right) = \mu_k E_t\left(\sigma^2_{t+j-2,k}\right) + \left(1-\mu_k\right) E_t\left(\sigma^2_{t+j-1}\right) \tag{2}$$

Define $\delta_{j,k} = E_t\left(\sigma^2_{t+j,k}\right)$ and $\gamma_j = E_t(\sigma^2_{t+j})$, the conditional equations are reduced to:

$$\gamma_j = \sum_{k=1}^{K} w_k \delta_{j,k} = \boldsymbol{w}'.\boldsymbol{\delta}_{j-1} \tag{3}$$

$$\delta_{j-1,k} = \mu_k \delta_{j-2,k} + \left(1-\mu_k\right)\gamma_{j-1} \tag{4}$$

where $\boldsymbol{w} = (w_1, w_2, \ldots, w_K)'$ and $\boldsymbol{\delta}_j = (\delta_{j,1}, \delta_{j,2}, \ldots, \delta_{j,K})'$. Plugging (3) into (4) produces:

$$\delta_{j-1,k} = \mu_k \delta_{j-2,k} + \left(1-\mu_k\right)\boldsymbol{w}'.\boldsymbol{\delta}_{j-2}. \tag{5}$$

In vector notation, (5) becomes:

$$\boldsymbol{\delta}_{j-1} = \boldsymbol{M}\boldsymbol{\delta}_{j-2} + \left(\boldsymbol{\iota}-\boldsymbol{\mu}\right)\boldsymbol{w}'.\boldsymbol{\delta}_{j-2} = \left[\boldsymbol{M} + \left(\boldsymbol{\iota}-\boldsymbol{\mu}\right)\boldsymbol{w}'\right]\boldsymbol{\delta}_{j-2} \tag{6}$$

where $\boldsymbol{\mu}$ is the vector of μ_k, $\boldsymbol{M}$ is the diagonal matrix consisting of μ_k, and $\boldsymbol{\iota}$ is the vector of one. (6) can be iterated *(j-1)* times:

$$\boldsymbol{\delta}_{j-1} = \left[\boldsymbol{M} + \left(\boldsymbol{\iota}-\boldsymbol{\mu}\right)\boldsymbol{w}'\right]^{j}\boldsymbol{\delta}_0 \tag{7}$$

where $\boldsymbol{\delta}_0 = (\sigma^2_{t,1}, \sigma^2_{t,2}, \ldots, \sigma^2_{t,K})'$ is the vector of EWMA volatilities at current time *t*. By combining (3) and (7), the conditional volatility at time *j* becomes:

$$E_t(\sigma^2_{t+j}) = \gamma_j = \boldsymbol{w}'.\boldsymbol{\delta}_{j-1} = \boldsymbol{w}'\left[\boldsymbol{M} + \left(\boldsymbol{\iota}-\boldsymbol{\mu}\right)\boldsymbol{w}'\right]^{j-1}\boldsymbol{\delta}_0. \tag{8}$$

Define $\boldsymbol{w'}_j = \boldsymbol{w}'\left[\boldsymbol{M} + \left(\boldsymbol{\iota}-\boldsymbol{\mu}\right)\boldsymbol{w}'\right]^{j}$ (hence $\boldsymbol{w}_0 = \boldsymbol{w}$), then $E_t(\sigma^2_{t+j}) = \boldsymbol{w'}_{j-1}.\boldsymbol{\delta}_0$. With serially uncorrelated returns, the optimal forecasts for the *h*-steps volatility (i.e., volatility over *h* steps), given information at time *t*, may be expressed as:

$$\sigma^2_{t+1:t+h} = \sum_{j=1}^{h} \sigma^2_{t+j} = \sum_{j=1}^{h} \boldsymbol{w'}_{j-1}.\boldsymbol{\delta}_0 \tag{9}$$

with $\boldsymbol{w}_j$ and $\boldsymbol{\delta}_0$ defined above.

The volatility forecast can be expressed in an alternative way. The one-step-ahead forecast of volatility is already given by

$$\sigma_{t+1}^2 = \sum_{k=1}^{K} w_k \sigma_{t,k}^2 = \sum_{k=1}^{K} w_k \left(1-\mu_k\right) \sum_{i=0}^{\infty} \mu_k^i r_{t-i}^2. \tag{10}$$

In practice, the sum over lags needs be cut off at some time T, and (10) becomes:

$$\sigma_{t+1}^2 = \sum_{k=1}^{K} w_k \left(1-\mu_k\right) \sum_{i=0}^{T} \mu_k^i r_{t-i}^2 = \sum_{i=0}^{T} \sum_{k=1}^{K} w_k \frac{\left(1-\mu_k\right)\mu_k^i}{1-\mu_k^T} r_{t-i}^2 = \sum_{i=0}^{T} \lambda(i) r_{t-i}^2 \tag{11}$$

with $\lambda(i) = \sum_{k=1}^{K} w_k \frac{\left(1-\mu_k\right)\mu_k^i}{1-\mu_k^T}$.

Applying the same recursive substitution procedure as above leads to the forecast for the *h*-step ahead volatility:

$$\sigma_{t+1:t+h}^2 = h \sum_{i=0}^{T} \lambda(h,i) r_{t-i}^2 \tag{12}$$

with the weights $\lambda\left(h,i\right)$ given by:

$$\lambda\left(h,i\right) = \sum_{k=1}^{K} \frac{1}{h} \sum_{j=1}^{h-1} w_{j,k} \frac{\left(1-\mu_k\right)}{1-\mu_k^T} \mu_k^i, \tag{13}$$

where $w_{j,k}$ is the *k* element of vector $\boldsymbol{w}_j$, and $\sum \lambda(h,i) = 1$. When $K = 1$, then $w = 1$, so the LM EWMA process reduces to an EWMA process with the forecast weights $\lambda\left(h,i\right) = \left(1-\mu_k\right)\mu_k^i \big/ \left(1-\mu_k^T\right)$, independent of forecast horizons.

Covariance Matrix Forecast

In a multivariate setting, the forecast of the *h*-steps covariance matrix, given the information set $\mathcal{F}_t$ at time *t*, is easily obtained:

$$\mathbf{H}_{t+1:t+h} = h \sum_{i=0}^{T} \lambda(h,i) \boldsymbol{r}_{t-i} \boldsymbol{r}_{t-i}^{'} \tag{14}$$

with $\lambda(h,i)$ defined in (13).

Appendix 5.2. Bayesian Prior Probabilities

Following Engle and Colacito (2006), I calculate the sample mean of non-overlapping consecutive subsamples of 63 days (3 months) from the full datasets. I discard all pairs of sample means that have at least one negative element. For the remaining *N* pairs, I compute $\theta_n = \frac{2}{\pi} arctan\left(\frac{\mu_{1,n}}{\mu_{2,n}}\right)$, solving the system:

$$\begin{cases} \mu_{1,n} = k\sin\left(\frac{\pi}{2}\theta_n\right) \\ \mu_{2,n} = k\cos\left(\frac{\pi}{2}\theta_n\right) \end{cases} \tag{1}$$

with $\forall n = 1,\ldots,N$. I use these values of θ to find parameters $\hat{a}$ and $\hat{b}$ that maximise the log-likelihood function of a Beta distribution:

$$\begin{aligned} \left(\hat{a},\hat{b}\right) &= \text{argmaxlogL}\left(\theta_1,\ldots,\theta_N; a,b\right) \\ &= \arg\max N\log\left(\frac{1}{\int_0^1 t^{a-1}\left(1-t\right)^{b-1}dt}\right) + \left(a-1\right)\sum_{i=1}^{N} log\left(\theta_i\right) + \left(b-1\right)\sum_{i=1}^{N} log\left(1-\theta_i\right). \end{aligned} \tag{2}$$

Finally, using the MLE $\hat{a}$ and $\hat{b}$, I compute the prior probability associated to each of the *N* pairs of sub sample mean:

$$Pr\left(\theta = \theta_i\right) = \frac{1}{\Upsilon}\frac{\theta_i^{\hat{a}-1}\left(1-\theta_i\right)^{\hat{b}-1}}{\int_0^1 t^{\hat{a}-1}\left(1-t\right)^{\hat{b}-1}dt}, \tag{3}$$

where Υ is the normalization factor.

Appendix 5.3. Comparison of Out-of-Sample Volatilities: Hedging International Portfolios

The table reports the average out-of-sample volatilities of the minimum variance hedging portfolios subject to the target excess return of 1. Each row reports the results of the hedging portfolio, in which the asset of the country in the first column are hedged against all other assets. The lowest volatility is normalised to 100.

	EWMA	GARCH DCC	LM EWMA	LM EWMA DCC	CGARCH DCC	Const
Panel A. International Stocks						
Australia	108.006	115.295	105.052	100.000	115.070	112.962
Austria	115.576	111.396	111.840	100.000	113.274	124.707
Belgium	120.686	101.276	117.778	100.000	103.684	111.501
Canada	112.535	106.420	109.291	100.000	116.915	110.464
Denmark	115.953	102.232	110.931	100.000	101.079	102.508
France	117.043	102.417	113.347	100.000	100.764	112.267
Germany	114.256	106.334	110.494	100.000	103.534	111.891
Hongkong	116.476	103.445	111.655	100.000	102.728	104.460
Ireland	125.153	105.140	125.247	103.454	111.355	100.000
Italy	111.064	100.500	105.284	100.000	100.852	102.242
Japan	117.539	100.950	112.753	100.000	102.605	117.311
Mexico	110.862	100.931	108.403	100.000	101.368	103.197
Netherland	112.997	105.749	108.559	100.000	106.228	126.765
New Zealand	114.199	104.664	110.012	100.000	103.651	112.213
Norway	112.849	102.394	108.409	100.000	104.159	122.970
Singapore	110.632	103.022	107.479	100.000	102.572	109.419
Spain	118.895	102.951	115.041	100.000	101.921	117.251
Sweden	111.627	108.682	107.583	100.000	107.923	111.715
Switzerland	107.679	105.526	104.179	100.795	100.000	109.504
UK	108.451	110.802	104.644	100.000	102.814	112.954
US	114.214	111.795	109.038	100.000	112.003	108.155
Panel B. International Bonds						
Austria	108.093	102.731	105.513	100.405	100.000	126.353
Belgium	112.533	103.177	110.531	100.000	102.654	114.578
Canada	115.398	105.282	111.594	100.000	106.334	106.435
Denmark	119.843	104.488	116.065	100.000	103.516	108.940
France	105.879	104.355	101.960	100.000	103.538	135.220
Germany	112.595	102.959	110.874	100.000	103.579	121.886
Ireland	112.693	100.026	108.723	100.000	101.489	120.162
Japan	112.614	101.036	108.469	100.000	100.318	102.660
Netherland	105.123	101.708	102.801	100.000	100.410	123.770
Sweden	107.576	100.374	104.929	100.000	100.144	105.054
Switzerland	110.430	101.447	106.907	100.793	100.000	112.192
UK	115.539	101.601	109.936	100.412	100.000	120.398
US	117.727	102.666	114.769	100.000	101.518	117.143

Appendix 5.4. Comparison of Out-of-Sample Volatilities: Hedging DJIA Portfolios

The table reports the average out-of-sample volatilities of the minimum variance hedging portfolios subject to the target excess return of 1. Each row reports the results of the hedging portfolio, in which the asset of the country in the first column are hedged against all other assets. The lowest volatility is normalised to 100.

Stocks	EWMA	GARCH DCC	LM EWMA	LM EWMA DCC	CGARCH DCC	Const
AA	137.637	102.048	113.701	100.000	102.707	108.131
AXP	137.160	101.638	113.252	100.000	101.546	106.162
BA	139.344	100.972	114.311	100.000	101.749	104.355
BAC	143.802	106.622	117.309	100.000	106.855	117.708
CAT	136.629	100.163	111.761	100.000	100.774	101.229
C	136.376	100.590	112.883	100.000	100.893	114.881
CVX	138.110	101.697	113.226	100.000	100.408	107.513
DD	135.070	100.897	111.231	100.000	100.459	113.425
DIS	133.373	100.000	110.310	100.312	100.493	103.552
GE	132.915	103.678	110.108	100.000	101.306	108.189
GM	139.040	105.369	113.505	100.000	103.228	112.008
HD	136.514	100.000	110.987	100.312	100.212	103.942
HPQ	138.465	102.360	113.865	102.144	102.906	100.000
IBM	139.064	100.403	113.156	100.000	100.550	100.805
INTC	136.395	101.552	112.602	100.611	101.197	100.000
JNJ	138.025	101.980	113.887	100.000	103.672	123.411
JPM	137.035	100.059	113.795	100.000	100.801	105.725
KO	138.825	100.217	114.373	100.000	100.607	103.023
MCD	136.311	100.468	111.773	100.164	101.264	100.000
MMM	136.469	102.391	113.026	102.531	103.237	100.000
MRK	136.509	101.063	112.848	100.000	101.383	108.407
MSFT	136.044	100.789	111.977	100.000	101.127	102.527
PFE	137.510	103.258	115.349	104.560	103.906	100.000
PG	140.641	101.110	115.274	100.000	101.275	101.859
T	141.839	101.032	115.291	100.000	101.599	102.881
UTX	132.020	100.000	109.834	100.238	102.110	103.287
VZ	135.632	101.360	111.542	100.000	100.897	106.616
WMT	137.245	100.414	113.353	100.000	100.677	104.799
XOM	135.983	100.971	112.014	100.000	100.545	109.576

Appendix 5.5. MAE for Longer Horizon Forecasts: Bivariate Systems

The table reports the MAE for each element of the forecast conditional covariance matrix over the forecast period. The benchmarks are the realised variances and covariances, proxied by the sum of squares and cross products of returns over the forecast horizons, respectively

	EWMA	GARCH DCC	LM EWMA	LM-EWMA DCC	CGARCH DCC
Panel A. One Week (5-Step) ahead Forecasts					
Variances					
Stock	4.830	4.788	4.809	4.809	4.818
Bond	0.524	0.555	0.512	0.512	0.566
S&P500	4.188	4.202	4.230	4.230	4.174
DJIA	3.960	3.984	3.987	3.987	3.973
Covariances					
Stock-Bond	1.008	1.023	0.990	0.992	1.041
S&P500-DJIA	3.974	3.971	4.023	4.026	3.971
Panel B. One Month (21-Step) ahead Forecasts					
Variances					
Stock	18.990	18.441	18.534	18.534	19.577
Bond	1.487	1.572	1.410	1.410	1.666
S&P500	16.624	16.457	16.299	16.299	16.798
DJIA	16.061	15.773	15.690	15.690	16.226
Covariances					
Stock-Bond	2.655	2.772	2.692	2.752	2.755
S&P500-DJIA	16.041	15.647	15.719	15.711	16.145
Panel C. One Quarter (63-Step) ahead Forecasts					
Variances					
Stock	54.875	62.730	55.299	55.299	64.307
Bond	3.769	3.983	3.547	3.547	4.331
S&P500	48.644	54.134	48.789	48.789	55.191
DJIA	46.157	50.160	46.426	46.426	49.948
Covariances					
Stock-Bond	5.801	7.186	5.864	7.006	6.882
S&P500-DJIA	46.498	50.509	46.615	46.462	50.777

Appendix 5.6. Comparison of Out-of-Sample Volatilities: Bivariate Portfolios with Monthly Rebalancing Frequency

The table reports the out-of-sample volatilities for the monthly rebalanced bivariate portfolios, constructed with the objective of minimizing variance subject to the target excess return of 1. Each row in the table reports the results for the pair of expected returns of the corresponding first two columns. The overall returns are the pair of weighted returns using the Bayesian prior probabilities. The lowest volatility in each row is normalised to 100.

Panel A. Stock-Bond Portfolio

μ_{Stock}	μ_{Bond}	EWMA	GARCH DCC	LM EWMA	LM EWMA DCC	CGARCH DCC	Const
0.00	1.00	100.000	100.120	100.360	100.000	100.240	103.121
0.16	0.99	100.236	100.118	100.472	100.000	100.236	102.243
0.31	0.95	100.455	100.114	100.682	100.000	100.341	101.477
0.45	0.89	100.642	100.214	100.856	100.000	100.535	100.642
0.59	0.81	100.689	100.295	100.984	100.098	100.689	100.000
0.71	0.71	100.799	100.622	101.066	100.355	101.066	100.000
0.81	0.59	100.078	100.312	100.312	100.000	100.703	100.156
0.89	0.45	100.000	100.678	100.203	100.339	100.949	101.762
0.95	0.31	100.000	100.858	100.286	100.458	101.430	102.460
0.99	0.16	100.000	101.495	100.386	100.241	102.943	101.447
1.00	0.00	100.000	106.548	101.592	103.183	103.183	103.229
Overall (weighted)		*100.000*	*100.907*	*100.291*	*100.485*	*101.357*	*102.534*

Panel B. S&P500-DJIA Portfolio

μ_{SP500}	μ_{DJIA}	EWMA	GARCH DCC	LM EWMA	LM EWMA DCC	CGARCH DCC	Const
0.00	1.00	100.704	105.352	100.986	100.141	105.070	100.000
0.16	0.99	100.595	104.875	100.951	100.000	104.637	100.119
0.31	0.95	100.665	104.183	101.046	100.000	104.278	100.570
0.45	0.89	100.635	102.680	100.987	100.000	103.173	102.186
0.59	0.81	100.476	100.000	100.571	100.143	100.905	105.619
0.71	0.71	101.062	105.178	100.000	101.460	107.600	102.456
0.81	0.59	100.865	100.956	100.000	100.319	102.641	107.696
0.89	0.45	100.417	105.003	100.000	100.695	105.281	106.741
0.95	0.31	100.094	104.143	100.000	100.847	105.367	104.049
0.99	0.16	100.000	104.260	100.237	100.947	105.444	102.840
1.00	0.00	100.000	104.628	100.281	100.842	105.470	102.104
Overall (weighted)		*100.933*	*104.828*	*100.000*	*101.353*	*107.264*	*102.743*

Appendix 5.7. Diebold–Mariano Joint Tests: Hedging International Stock and Bond Portfolios with Different Rebalancing Frequencies

The table reports the *t*-statistics of the Diebold–Mariano joint tests for the hedging international stock and bond portfolios, using the improved test of Engle and Colacito (2006). The *t*-statistics for the standard test are reported in parentheses. A positive number indicates that the model in the row is better than the model in the column, and vice-versa. *, ** and *** denote rejection of the equality hypothesis of the two models at 10%, 5% and 1% significance level.

	EWMA	GARCH DCC	LM-EWMA	LM-EWMA DCC	CGARCH DCC
Panel A. Monthly Rebalancing					
EWMA	–	-1.66* (-2.62***)	-2.40** (-8.03***)	-2.56** (-7.53***)	-2.91*** (-6.57***)
GARCH-DCC	1.66* (2.62***)	–	1.16 (1.00)	-2.68*** (-3.26***)	-2.02** (-2.22**)
LM-EWMA	2.40** (8.03***)	-1.16 (-1.00)	–	-1.99** (-5.33***)	-2.12** (-4.30***)
LM-EWMA-DCC	2.56** (7.53***)	2.68*** (3.26***)	1.99** (5.33***)	–	1.60 (3.05***)
CGARCH-DCC	2.91*** (6.57***)	2.02** (2.22**)	2.12** (4.30***)	-1.60 (-3.05***)	–
Panel B. Quarterly Rebalancing					
EWMA	–	-1.48 (-3.01***)	-2.38** (-5.65***)	-1.61 (-4.15***)	-1.73* (-3.77***)
GARCH-DCC	1.48 (3.01***)	–	1.18 (1.64)	-1.91* (-3.33***)	-1.80* (-1.70*)
LM-EWMA	2.38** (5.65***)	-1.18 (-1.64)	–	-1.42 (-3.17***)	-1.48 (-2.61***)
LM-EWMA-DCC	1.61 (4.15***)	1.91* (3.33***)	1.42 (3.17***)	–	0.30 (2.54**)
CGARCH-DCC	1.73* (3.77***)	1.80* (1.70*)	1.48 (2.61***)	-0.30 (-2.54**)	–

Appendix 5.8. Comparison of Out-of-Sample Volatilities: Bivariate Portfolios with Different Estimation Windows and Rebalancing Frequencies

The table reports the out-of-sample volatilities for the bivariate portfolios, constructed with the objective of minimizing variance subject to the target excess return of 1, across different estimation windows and rebalancing frequencies. The portfolios are constructed based on the overall Bayesian prior weighted returns and the conditional covariance matrices generated from different conditional volatility models. The FIGARCH-DCC model is excluded when the short 2-year estimation window is applied. Conditional volatilities are reported in parentheses. The lowest volatility in each row is normalised to 100.

	EWMA	GARCH DCC	LM EWMA	LM EWMA DCC	CGARCH DCC	FIGARCH DCC
Panel A. 2-Year Estimation Window						
Panel A1. Stock-Bond Portfolio						
Daily	100.30 (100.10)	101.58 (101.81)	100.00 (100.00)	101.10 (100.81)	103.01 (103.27)	–
Weekly	100.18 (100.47)	100.00 (100.00)	100.07 (100.25)	100.25 (100.73)	101.17 (101.38)	–
Monthly	100.00 (100.00)	100.01 (101.67)	100.29 (100.20)	100.84 (100.59)	102.36 (103.19)	–
Quarterly	100.11 (100.00)	101.20 (101.95)	100.27 (101.14)	100.00 (100.75)	101.18 (102.74)	–
Panel A2. S&P500-DJIA Portfolio						
Daily	100.10 (100.23)	101.36 (101.50)	100.00 (100.00)	100.48 (100.45)	102.06 (102.16)	–
Weekly	100.00 (100.00)	102.25 (102.01)	100.67 (100.88)	100.91 (101.04)	102.45 (102.62)	–
Monthly	101.27 (101.21)	100.47 (101.24)	100.00 (100.31)	100.33 (100.00)	101.93 (102.49)	–
Quarterly	101.36 (101.87)	101.31 (102.17)	100.00 (100.00)	101.45 (100.98)	102.42 (103.61)	–
Panel B. 5-Year Estimation Window						
Panel B1. Stock-Bond Portfolio						
Daily	100.21 (100.04)	102.01 (101.95)	100.00 (100.00)	101.30 (100.93)	103.45 (103.34)	115.11 (114.07)
Weekly	100.93 (101.39)	101.91 (102.58)	100.00 (100.00)	102.15 (102.09)	100.08 (101.75)	116.19 (115.06)
Monthly	100.00 (100.00)	103.42 (102.56)	100.36 (100.46)	102.71 (101.41)	100.65 (103.34)	111.09 (105.35)
Quarterly	100.00 (100.00)	105.15 (107.11)	101.13 (102.06)	102.98 (103.16)	106.73 (106.53)	125.05 (124.26)
Panel B2. S&P500-DJIA Portfolio						
Daily	100.12 (100.27)	101.21 (101.13)	100.00 (100.00)	100.45 (100.22)	101.40 (101.33)	112.27 (108.58)
Weekly	100.57 (100.52)	100.81 (100.73)	101.25 (100.52)	100.00 (100.00)	102.99 (101.94)	109.95 (108.01)
	101.45	101.61	100.00	100.33	104.32	102.71

	EWMA	GARCH DCC	LM EWMA	LM EWMA DCC	CGARCH DCC	FIGARCH DCC
Monthly	(101.71)	(102.28)	(100.77)	(100.00)	(105.09)	(103.14)
Quarterly	102.91	100.00	101.19	102.65	105.67	102.95
	(103.94)	(100.00)	(101.98)	(102.29)	(107.23)	(101.94)

Panel C. 10-Year Estimation Window

Panel C1. Stock-Bond Portfolio

Daily	100.19	100.65	100.00	101.07	102.92	117.64
	(100.17)	(101.06)	(100.00)	(100.77)	(103.58)	(116.41)
Weekly	100.26	100.38	100.00	101.64	100.13	115.68
	(100.00)	(100.97)	(100.05)	(101.56)	(101.50)	(117.57)
Monthly	101.18	100.26	101.45	102.26	100.00	108.70
	(100.45)	(100.08)	(100.00)	(100.57)	(100.11)	(103.67)
Quarterly	100.23	101.69	100.00	100.53	100.33	129.08
	(100.79)	(101.53)	(101.67)	(100.81)	(100.00)	(126.37)

Panel C2. S&P500-DJIA Portfolio

Daily	100.11	100.50	100.00	100.29	101.14	117.46
	(100.28)	(100.47)	(100.00)	(100.20)	(101.40)	(112.55)
Weekly	100.98	100.00	101.60	100.21	101.67	112.04
	(100.22)	(100.00)	(100.50)	(100.05)	(100.59)	(109.38)
Monthly	101.71	100.74	100.00	100.30	101.88	107.28
	(101.70)	(101.38)	(100.41)	(100.14)	(100.00)	(106.84)
Quarterly	103.66	101.68	101.89	103.28	100.00	114.98
	(105.83)	(102.68)	(103.45)	(103.96)	(100.00)	(110.17)

Appendix 5.9. Diebold–Mariano Joint Tests: Hedging DJIA Portfolios with Different Estimation Windows

The table reports the *t*-statistics of the Diebold–Mariano joint tests for the hedging DJIA portfolios, using the improved test of Engle and Colacito (2006). The *t*-statistics for the standard test are reported in parentheses. A positive number indicates that the model in the row is better than the model in the column, and vice-versa. *, ** and *** denote rejection of the equality hypothesis of the two models at 10%, 5% and 1% significance level.

Panel A. 2-Year Estimation Window

		EWMA	GARCH DCC	LM EWMA	LM EWMA DCC	CGARCH DCC
Daily	LM EWMA	10.1*** (40.2***)	-5.44*** (-28.2***)	–	-8.95*** (-27.4***)	-1.33 (-26.2***)
	LM EWMA DCC	9.96*** (36.4***)	1.86* (3.47***)	8.95*** (27.4***)	–	1.17 (2.71)
Weekly	LM EWMA	6.60*** (18.8***)	-6.58*** (-17.7***)	–	-6.03*** (-16.7***)	-7.41*** (-18.0***)
	LM EWMA DCC	7.33*** (20.4***)	1.79* (1.67*)	6.03*** (16.7***)	–	1.36 (2.20**)
Monthly	LM EWMA	7.33*** (15.88***)	-4.03*** (-9.76***)	–	-5.18*** (-9.37***)	-5.00*** (-9.73***)
	LM EWMA DCC	6.24*** (13.11***)	-0.55 (0.36)	5.18*** (9.37***)	–	0.34 (1.20)
Quarterly	LM EWMA	6.73*** (12.17***)	0.39 (-4.11***)	–	-2.62*** (-5.47***)	-3.86*** (-6.00***)
	LM EWMA DCC	4.89*** (9.50***)	1.08 (0.61)	2.62*** (5.47***)	–	-1.39 (-0.31)

Panel B. 10-Year Estimation Window

		EWMA	GARCH DCC	LM EWMA	LM EWMA DCC	CGARCH DCC	FIGARCH DCC
Daily	LM EWMA	7.56*** (27.2***)	-6.15*** (-18.6***)	–	-6.38*** (-18.3***)	-6.36*** (-18.9***)	2.82** (3.21***)
	LM EWMA DCC	6.89*** (24.5***)	2.44** (3.79***)	6.38*** (18.3***)	–	2.04** (3.94***)	3.35*** (7.29***)
Weekly	LM EWMA	5.75*** (15.4***)	-4.38*** (-12.9***)	–	-4.73*** (-12.9***)	-5.39*** (-13.7***)	2.24** (-0.85)
	LM EWMA DCC	5.36*** (15.8***)	2.49** (3.28***)	4.73*** (12.9***)	–	1.86* (2.83***)	3.28*** (6.06***)
Monthly	LM EWMA	6.17*** (13.2***)	-3.26*** (-7.01***)	–	-4.00*** (-7.55***)	-4.21*** (-7.15***)	0.47 (-2.17**)
	LM EWMA DCC	4.74*** (10.6***)	-0.23 (1.30)	4.00*** (7.55***)	–	1.33 (1.88*)	1.96** (2.28**)
Quarterly	LM EWMA	4.79*** (8.19***)	-3.29*** (-5.33***)	–	-4.18*** (-6.31***)	-0.61 (-2.92***)	1.07 (1.14)
	LM EWMA DCC	4.63*** (8.02***)	-0.36 (1.39)	4.18*** (6.31***)	–	1.93* (2.29**)	1.11 (1.47)

Appendix 5.10. Diebold–Mariano Joint Tests: Hedging International Stock and Bond Portfolios with 10-Year Estimation Window

The table reports the *t*-statistics of the Diebold–Mariano joint tests for the hedging international stock and bond portfolios, using the improved test of Engle and Colacito (2006). The *t*-statistics for the standard test are reported in parentheses. A positive number indicates that the model in the row is better than the model in the column, and vice-versa. *, ** and *** denote rejection of the equality hypothesis of the two models at 10%, 5% and 1% significance level.

	EWMA	GARCH DCC	LM-EWMA	LM-EWMA DCC	CGARCH DCC
Panel A. Weekly Rebalancing					
EWMA	–	-3.547*** (-8.377***)	-4.461*** (-9.362***)	-3.981*** (-9.064***)	2.949*** (4.496***)
GARCH-DCC	3.547*** (8.377***)	–	3.081*** (6.257***)	-2.560** (-4.825***)	2.836*** (5.681***)
LM-EWMA	4.461*** (9.362***)	-3.081*** (-6.257***)	–	-3.934*** (-8.045***)	3.170*** (5.064***)
LM-EWMA-DCC	3.981*** (9.064***)	2.560** (4.825***)	3.934*** (8.045***)	–	2.960*** (5.758***)
CGARCH-DCC	-2.949*** (-4.496***)	-2.836*** (-5.681***)	-3.170*** (-5.064***)	-2.960*** (-5.758***)	–
Panel B. Monthly Rebalancing					
EWMA	–	-3.934*** (-6.108***)	-4.402*** (-8.584***)	-5.211*** (-8.969***)	1.868* (3.770***)
GARCH-DCC	3.934*** (6.108***)	–	1.512 (2.425**)	-2.721*** (-3.686***)	1.929* (4.473***)
LM-EWMA	4.402*** (8.584***)	-1.512 (-2.425**)	–	-4.049*** (-5.665***)	2.000** (4.148***)
LM-EWMA-DCC	5.211*** (8.969***)	2.721*** (3.686***)	4.049*** (5.665***)	–	2.054** (4.602***)
CGARCH-DCC	-1.868* (-3.770***)	-1.929* (-4.473***)	-2.000** (-4.148***)	-2.054** (-4.602***)	–
Panel C. Quarterly Rebalancing					
EWMA	–	-0.010 (-2.014**)	-3.648*** (-7.071***)	-1.470 (-4.371***)	2.383** (2.762***)
GARCH-DCC	0.010 (2.014**)	–	-0.640 (-0.264)	-1.438 (-2.466**)	2.639*** (3.238***)
LM-EWMA	3.648*** (7.071***)	0.640 (0.264)	–	-0.384 (-1.679*)	2.568*** (3.208***)
LM-EWMA-DCC	1.470 (4.371***)	1.438 (2.466**)	0.384 (1.679*)	–	2.732*** (3.416***)
CGARCH-DCC	-2.383** (-2.762***)	-2.639*** (-3.238***)	-2.568*** (-3.208***)	-2.732*** (-3.416***)	–

Chapter 6

The Economic Value of Long Memory Volatility Timing

Extensive research suggests significant economic benefits to exploiting forecasts of multivariate conditional volatility models relative to using the unconditional covariance matrix estimator. Exploiting the predictability of volatility and covariance has become a key driver in many applied areas of finance, including asset allocation, asset pricing and risk management. Fleming et al. (2001) are among the first to study the economic value of predicting and timing volatility for risk averse investors in an asset allocation setting. Specifically, they show that investors are better off in terms of utility when switching from a static to a dynamic volatility timing strategy. Recent studies incorporate more properties of volatility dynamics in application to investment decisions. Thorpe and Milunovich (2007) allow for asymmetries in modelling volatility and correlation, and show that investors are willing to pay to switch from symmetric to asymmetric forecasts. Similarly, Hyde et al. (2010) demonstrate the benefits of accounting for volatility jumps in asset allocation strategies. In these dynamic economic value studies, the conditional covariance matrix is typically estimated applying popular conditional volatility models such as the multivariate EWMA or multivariate GARCH models, where shocks to volatility and covariance dissipate rapidly due to their exponential weighting. Consequently, most of the studies on the economic value of the short memory conditional covariance matrix focus on short horizon day traders. While this approach may make the most use of the forecast power of the short memory conditional volatility models, it may not nevertheless correspond to the needs of most practical investors, who often rebalance their portfolios at lower frequencies.

This chapter examines the economic value of allowing for long memory volatility dynamics in forecasting the covariance matrix for asset allocation over both short and long horizons. In Chapter 5, multivariate long memory conditional volatility models have been found to produce better forecasts of the covariance matrix than those produced by multivariate short memory volatility models, especially for long horizons. The four multivariate long memory volatility models (the LM-EWMA, LM-EWMA-DCC, FIGARCH-DCC and CGARCH-DCC models) are now compared with the two multivariate short memory volatility models (the multivariate Riskmetrics EWMA and GARCH-DCC models) in terms of economic performance using the volatility timing framework of Fleming et al. (2001). For simplicity, I concentrate primarily on the one-period mean-variance portfolio choice and ignore the hedging demands caused by time-varying investment opportunities. The four data systems described in Chapter 4 are again employed to construct the optimal portfolios. Expected returns are assumed constant and investors periodically update their portfolios based on forecasts of the covariance matrix generated from conditional volatility models, i.e., they predict and time volatility. Dynamic portfolios constructed with alternative conditional volatility models are also evaluated against static portfolios constructed with the constant unconditional covariance matrix estimates, and equally-weighted portfolios. Portfolio performance is evaluated using the out-of-sample Sharpe ratio, the abnormal return and the performance fee that investors are willing to pay to switch from the static to the dynamic strategies. The effects of transaction costs are also considered. The research reports three main findings. First, consistent with the literature, the dynamic volatility timing strategies significantly outperform the static strategies in terms of different performance measures and across different rebalancing frequencies. Second, incorporating long memory in volatility dynamics brings further economic gains. The long memory volatility models consistently produce portfolios that are more economically useful than those produced by the short memory volatility models at all investment horizons. Among the long memory models, the two parsimonious LM-EWMA models generally dominate. Third, when transaction costs are taken into account, the gains from daily rebalanced dynamic portfolios deteriorate. However, it is still worth implementing the dynamic strategies at lower rebalancing frequencies. The results apply to all four datasets, both low and high dimensions, and are robust to estimation error in expected returns, the choice of risk aversion coefficient and the use of a long-only constraint.

The remainder of the chapter is structured as follows. Section 6.1 introduces the dynamic volatility timing framework applied to evaluate the economic benefits of the multivariate conditional volatility models. The empirical results are presented in Section 6.2, while Section 6.3 offers some concluding comments and suggestions for future research.

6.1 The Economic Value of Dynamic Volatility Timing Strategy

6.1.1 The Dynamic Volatility Timing Framework

Suppose that an investor allocates $\mathbf{w}_t$ fraction of his wealth to n risky assets and the remainder $\left(1-\mathbf{w}_t^{'}\boldsymbol{\iota}\right)$ to a risk-free asset, where $\boldsymbol{\iota}$ is the $n\times1$ unit vector. Given the mean-variance optimization framework, the investor maximises his expected utility U_{t+1}:

$$\max_{\mathbf{w}_t}\left\{E\left(U_{t+1}\right)=\mu_{p,t+1}-\frac{\lambda}{2}\sigma_{p,t+1}^{2}\right\} \tag{6.1}$$

where $\mu_{p,t+1}$ is the portfolio's expected returns $\mu_{p,t+1}=\mathbf{w}_t^{'}\boldsymbol{\mu}_{t+1}+\left(1-\mathbf{w}_t^{'}\boldsymbol{\iota}\right)r^f$, $\sigma_{p,t+1}^{2}$ is the portfolio's expected variance $\sigma_{p,t+1}^{2}=\mathbf{w}_t^{'}\mathbf{H}_{t+1}\mathbf{w}_t$, $\boldsymbol{\mu}_{t+1}$ is the vector of expected returns, $\mathbf{H}_{t+1}$ is the conditional covariance matrix, r^f is the risk-free rate and λ is the risk aversion coefficient. Following Fleming et al. (2001), I assume constant expected returns $\boldsymbol{\mu}_{t+1}\equiv\boldsymbol{\mu}$ so as to specifically examine the economic value of volatility timing. In the empirical study, I assume a risk free rate of 4% and a risk aversion coefficient of 1. Different values of λ are later considered in the robustness test. Short sales are allowed and no transaction costs are included. The solution to this optimization problem is:

$$\mathbf{w}_t^{*}=\frac{1}{\lambda}\mathbf{H}_{t+1}^{-1}\left(\boldsymbol{\mu}-\boldsymbol{\iota}r^f\right). \tag{6.2}$$

If the covariance matrix is constant, the optimal weights will be constant over time, which is referred to as the '*static strategy*'. However, if the investor believes that the covariance matrix is time-varying, he will follow a '*dynamic strategy*' to change the optimal weights based on his forecasts of the conditional covariance matrix. The investor will employ the six multivariate conditional volatility models studied in Chapter 5 to generate forecasts of the covariance matrix for the dynamic strategies. By comparing the performance of the static and dynamic portfolios, I can evaluate the economic value of volatility timing. The portfolios constructed with the four multivariate long memory volatility models (the LM-EWMA, LM-EWMA-DCC, FIGARCH(1,*d*,1)-DCC and CGARCH(1,1)-DCC models) are also compared to those constructed with the two short memory volatility models (the multivariate EWMA and the GARCH(1,1)-DCC models) to specifically evaluate the gains of exploiting long memory vs. short memory properties of volatility.

6.1.2 Performance Measures of the Dynamic Strategies

The performance of the optimal portfolios is evaluated using three common performance measures. First, the out of sample Sharpe ratio of each portfolio is calculated as the sample mean of the realised portfolio excess returns over the risk free rate divided by their sample standard deviation, $SR = \left(\mu_p - r_f\right)/\sigma_p$. Though the Sharpe ratio is the most common portfolio performance measure, Han (2006) argues that the ex post Sharpe ratio may be misleading in the sense that it does not take into account the time-varying conditional volatility. Using the realised sample standard deviation, the ex post Sharpe ratio may overestimate the conditional risk that an investor faces in a dynamic strategy, hence underestimating the performance of this strategy. Therefore, I additionally consider a Sharpe-related measure that compares the two portfolio performance on the same risk footing. In particular, I consider the abnormal return measure *M2* of Modigliani and Modigliani (1997) that the dynamic strategy would earn if it had the same risk as the static benchmark.

$$M2 = \frac{\sigma_s}{\sigma_d}\left(\mu_d - r_f\right) - \left(\mu_s - r_f\right) \tag{6.3}$$

where μ_s, σ_s and μ_d, σ_d are the out-of-sample means and standard deviations of the static and dynamic portfolios, respectively. Note that this measure is directly related to the Sharpe ratios of the two strategies as $M2 = \sigma_s\left(SR_d - SR_s\right)$.

The third measure is the performance fee, suggested in Fleming et al. (2001) and now the most widely used performance measure in the volatility timing literature. Consider an investor with the quadratic utility function:

$$U\left(W_{t+1}\right) = W_{t+1} - \frac{a}{2}W_{t+1}^2, \tag{6.4}$$

where W_{t+1} is the investor's expected wealth. Assume that each day the investor fixes the amount of the initial wealth W_0. Fleming et al. (2001) also fix the coefficient of relative risk aversion $\gamma_t = \frac{aW_t}{1-aW_t}$ constant, and examine the average utility function:

$$\bar{U}(\cdot) = W_0 \left(\sum_{t=0}^{T-1} R_{p,t+1} - \frac{\gamma}{2(1+\gamma)} R_{t+1}^2 \right). \tag{6.5}$$

By fixing γ constant, the investor implicitly follows the approximation of a second order Taylor expansion of a non-quadratic utility. Constant relative risk aversion also implies that expected utility is linearly homogenous in wealth. The performance fee Δ is defined as the maximum fee that the investor would be willing to pay to switch from a static strategy to a dynamic strategy, without being worse off in terms of utility. To estimate this fee, I find the value of Δ that equates the realised average utilities for two alternative portfolios:

$$\sum_{t=0}^{T-1} \left(R_{d,t+1} - \Delta \right) - \frac{\gamma}{2(1+\gamma)} \left(R_{d,t+1} - \Delta \right)^2 = \sum_{t=0}^{T-1} R_{s,t+1} - \frac{\gamma}{2(1+\gamma)} R_{s,t+1}^2 , \tag{6.6}$$

where $R_{d,t+1}$ and $R_{s,t+1}$ are the gross realised returns of the dynamic and static strategies, respectively. In the empirical analysis, I report the annualised performance fees in basis points for two different values of constant relative risk aversion coefficients γ, 1 and 5.

6.1.3 Transaction Costs

Volatility timing requires regular updates of portfolios, thus incurring non-trivial transaction costs. Transaction costs may be high enough to offset all the gains from the dynamic strategies. Transaction costs are hence examined to see if it is still worth following the dynamic strategies. Following Han (2006), I estimate the breakeven transaction cost τ^{be}, defined as the transaction cost that make investors indifferent between the dynamic and the static strategies in terms of utility. If an investor has a transaction cost lower than the breakeven transaction cost, he will be better off with the dynamic strategy; otherwise he should follow the static benchmark. Han sets the transaction costs equal to a fixed percentage (τ) of the value traded for all stocks. The costs for the static and dynamic strategies are given by

$$\tau \left| w - \frac{w\left(1 + r_{p,t+1}\right)}{w\left(r_{p,t+1} - r_f\right) + r_f + 1} \right| \tag{6.7}$$

and $\tau\left|w_{t+1}-\frac{w_t\left(1+r_{p,t+1}\right)}{w_t\left(r_{p,t+1}-r_f\right)+r_f+1}\right|$, respectively. (6.8)

The breakeven transaction cost is computed by equating the utilities of the static and dynamic strategies after taking into account the trading costs. The higher the breakeven transaction cost, the more easily the dynamic trading strategies can be implemented. Since the breakeven transaction is a proportional cost paid every time the portfolios are rebalanced, I report this cost in basis points at the rebalancing frequency, e.g., for a daily rebalanced portfolio, I report the cost in daily basis points. Breakeven transaction costs are only estimated when the performance fees in (6.6) are positive.

6.2 Empirical Results

6.2.1 Low Dimensional Systems: The Stock-Bond and S&P500-DJIA Portfolios

For each portfolio, the whole sample is divided into an estimation period (1 January 1988 to 31 December 1993, 1517 daily observations) and a forecast period (1 January 1994 to 31 December 2009, 4031 observations). Expected returns are assumed to be constant and set to the sample mean of the estimation period. The investor actively rebalances his portfolios periodically, based on changes in forecasts of the conditional covariance matrix. The estimation period is used to initiate the estimation of the conditional volatility models and generate one-step-ahead forecasts of the covariance matrix. The forecasts are then used, along with the constant expected returns, to compute the optimal portfolio weights. Realised portfolio returns at the next step are calculated. Then the estimation window is rolled forward one step, models re-estimated, forecasts made, portfolios rebalanced and realised returns calculated, and so on until the end of the sample is reached. The realised performance of the dynamic portfolios will be compared with that of the ex ante optimal static portfolio, constructed based on the sample mean and covariance matrix of the estimation period. Another benchmark is the equally weighted portfolio.

Table 6.1 reports the out of sample performance of the daily rebalanced static and dynamic strategies. To estimate the FIGARCH model, I follow standard practice to use a truncation lag of 1000. It is clearly demonstrated that nearly all of the dynamic portfolios, except those constructed with the FIGARCH-DCC model, outperform the static and the equally weighted portfolios. Conditional volatility models consistently produce portfolios with higher Sharpe ratios and positive abnormal returns. The passive investor is also willing to pay annualised performance fees of 9 to 41 *bps* for the Stock-Bond portfolio and of 3 to 8 *bps* for the S&P500-DJIA portfolio to switch from the static to the dynamic strategies. These seemingly small fees are equal to 30% up to 150% of the realised average returns of the static strategies. These findings are consistent with the literature (see, for example, Fleming et al., 2001, 2003, Han, 2006), confirming the value of volatility timing in asset allocation. Note that an investor with a low relative risk aversion coefficient γ is more inclined to take risk, and vice versa. As a result, he is willing to pay more to switch to the more risky portfolios. For example, in the Stock-Bond system, where the static portfolio has lower risk than the dynamic portfolios, the performance fees that an investor with $\gamma = 1$ is willing to pay are higher than those for an investor with $\gamma = 5$. Conversely, if the static portfolio has higher risk (as in the S&P500-DJIA portfolio), the more risk averse investor with high γ is more willing to pay to switch to the dynamic strategies. Incorporating long memory in volatility dynamics brings further economic gains. The long memory LM-EWMA, LM-EWMA-DCC and CGARCH-DCC models generally produce portfolios with higher Sharpe ratios, abnormal returns and performance fees than the short memory GARCH-DCC model. The heavy parameterisation and computational burden evidently hinder the performance of the FIGARCH model, which systematically generates the least desirable portfolios. While the LM-EMWA-DCC model is the most economically useful model among the DCC models, embedding long memory volatility in the EWMA structure does not bring material benefits. The parsimony of the EWMA model offsets the possible gains from the more correctly specified, yet more complex LM-EWMA model. The EWMA model even outperforms the LM-EWMA model in terms of performance fees.

Practical portfolio management often requires longer investment horizons. Thus I evaluate the benefits of the dynamic volatility timing strategies for horizons of up to one month. Tables 6.2 and 6.3 compare the out of sample performance of the weekly and monthly rebalanced portfolios, respectively. As previously, the equally weighted portfolios perform worse than the mean-variance efficient portfolios. Despite longer horizons, the performance of the dynamic strategies remains strong, with even higher Sharpe ratios and higher performance fees in most cases. For example, the Sharpe ratios of the LM-EWMA and LM-EWMA-DCC portfolios nearly double those of the static portfolios. Consistent with the previous findings, the long memory volatility models significantly outperform the short memory GARCH-DCC model in most scenarios, again with the exception of the FIGARCH-DCC model. Among the long memory models, the LM-EWMA and LM-EWMA-DCC models generally dominate the CGARCH-DCC and FIGARCH-DCC models. The EWMA model still proves itself a simple, yet economically useful model. It is interesting to see the DCC models perform better in the high correlated S&P500-DJIA portfolio than in the low correlated Stock-Bond portfolio. Experiments with a quarterly investment horizon, though not reported here, yield similar results.

I now have a closer look on the gains of switching from the short memory GARCH-DCC model to the long memory volatility alternatives. Table 6.4 suggests that the investor benefits from better covariance matrix estimators in his investment decisions. He is willing to pay to switch not only from the static to the dynamic portfolios, but also from the short memory GARCH to the long memory volatility timing strategies. For example, the investor readily pays around 26 *bps* to switch from the daily rebalanced Stock-Bond GARCH-DCC portfolio to the corresponding LM-EWMA portfolio without being worse off in terms of utility. The advantage of a simple model is also clearly demonstrated when the two parsimonious LM-EWMA models dominate the more complicated CGARCH-DCC model. The FIGARCH-DCC model, again, produces poor results. Though the FIGARCH-DCC portfolio unexpectedly generates the highest performance fees in the monthly rebalanced Stock-Bond portfolio, this is more likely to be a result of luck than as proof of a robust model. However, experiments with the short memory EWMA model produce contradictory results (Table 6.5). Though the EWMA portfolio underperforms the long memory portfolios in some cases in terms of the Sharpe ratio and the abnormal fee, it dominates in terms of the performance fee. The investor now is willing to pay to switch from the long memory portfolios back to the short memory EWMA portfolio. The long memory models, though correctly specified, may have high estimation error, which makes them underperform the misspecified short memory EWMA yet with less estimation error. This demonstrates the trade-off between estimation error and specification error.

Dynamic strategies requires more trading than static strategies. In some cases, transaction costs may be high enough to offset all the gains from the dynamic strategies. To get a sense of the amount of trading required with each strategy, the breakeven transaction costs are calculated and reported in Table 6.6. Note that the higher the breakeven transaction costs, the easier the dynamic portfolios to be implemented. The breakeven transaction costs of the daily rebalanced dynamic portfolios are quite low, especially for the S&P500-DJIA portfolio, suggesting that volatility timing may not be desirable for a day trader. For example, the breakeven transaction costs of the LM-EWMA portfolio are around 9 *bps* for the Stock-Bond portfolio and only 2 *bps* for the S&P500-DJIA portfolio. With less frequent trading, the weekly and monthly rebalanced portfolios yield much higher breakeven transaction costs, making the dynamic strategies more feasible. The breakeven costs of the weekly and monthly rebalanced Stock-Bond portfolios constructed with the LM-EWMA model increase to 19 and 73 *bps*, respectively. Among the conditional volatility models, the EWMA, LM-EWMA and LM-EWMA-DCC models consistently produce portfolios that are not only the most superior in terms of Sharpe ratios, abnormal returns and performance fees, but also the most feasible in terms of transaction costs. The breakeven transaction costs of the LM-EWMA portfolio, for example, are much higher than those of the GARCH-DCC, CGARCH-DCC and FIGARCH-DCC portfolios in both datasets and across all rebalancing frequencies.

Estimation Error in Expected Returns

Constructing optimal portfolios requires estimation of expected returns. The results so far are based on the constant sample returns of the estimation period. Fleming et al. (2001) suggest that using one vector of expected returns may be inappropriate. In this section, I consider a range of expected returns to control for their estimation error. Again, I follow Engle and Colacito (2006) to use all possible pairs of relative expected returns in the form of $\mu = \left[\sin\frac{\pi j}{20}, \cos\frac{\pi j}{20}\right]$, for $j \in \{0,\ldots,10\}$. In the bivariate asset allocation framework, using the relative returns is sufficient to calculate the out of sample Sharpe ratio, which depends on the relative, not the absolute returns. Different absolute returns with the same proportion just move the optimal portfolio along the efficient frontier. This changes the portfolio's expected return and volatility, but not its Sharpe ratio. Also following Engle and Colacito (2006), I estimate a weighted vector of expected returns using the Bayesian prior probabilities.[10] The weighted vector of expected returns is higher than the sample mean since by construction, it does not include negative returns.

[10] Details of Engle and Colacito's (2006) weighted vector of expected returns have been described in Chapter 5.

Table 6.7 reports the out of sample Sharpe ratios of the optimal portfolios constructed using different covariance matrix estimators with all pairs of expected returns. It is clearly demonstrated that portfolios constructed with the pairs of expected returns that are close to the true vector of expected returns, such as the weighted returns, $[\mu_{Stock}, \mu_{Bond}] = [0.95, 0.31]$ and $[0.99, 0.16]$, and $[\mu_{S\&P500}, \mu_{DJIA}] = [0.59, 0.81]$ and $[0.71, 0.71]$, generally yield the highest Sharpe ratios. Portfolio performance is similar to that with the ex ante sample mean. The static portfolios are generally dominated by the dynamic portfolios, among which the long memory portfolios are consistently the most desirable. As previously, the long memory LM-EWMA and LM-EWMA-DCC models outperform the short memory GARCH-DCC model and other long memory CGARCH-DCC and FIGARCH-DCC models in most vectors of expected returns. In particular, the LM-EWMA model performs best in the Stock-Bond portfolio with the Sharpe ratio, for example, equal to 0.5625 using the weighted returns, while the LM-EWMA-DCC model dominates in the S&P500-DJIA portfolio, with a corresponding Sharpe ratio of 0.3227. The EWMA model, again, produces quite impressive portfolio performance.

I then estimate the performance fee that the investor is willing to pay to switch from the static to the dynamic strategies. Note that the Sharpe ratio does not depend on the absolute returns, but the performance fee does. The magnitude of the performance fee is affected by the magnitude of the vector of expected returns. So when comparing the performance fees of alternative portfolios, I am concerned not with the absolute gains, but with the relative gains of the dynamic strategies. Table 6.8 reports the relative performance fees with $\gamma = 1$. For each vector of expected return, the highest positive performance fee is normalised to 100. A positive number means that the investor is willing to pay to switch from the static to the dynamic strategies. A "–" sign indicates a negative performance fee, and so the investor will stay with the static strategy. A number of 95 means that the performance fee of the strategy in the corresponding column is equal to 95% of the highest performance fee. It is interesting to see that the simple EWMA model generally yields the highest performance fees. The two long memory LM-EWMA and LM-EWMA DCC models, though inferior to the EWMA model, consistently outperform the remaining models. The heavy parameterisation and computational burden evidently hinder the performance of the other two long memory models, especially the FIGARCH-DCC model. While the investor is still slightly better off with the CGARCH-DCC portfolio than with the short memory GARCH-DCC portfolio, he will not under any circumstances switch to the FIGARCH-DCC portfolio, even from the static one. Again, the DCC models perform better in the high correlated S&P500-DJIA portfolio. An experiment with $\gamma = 5$ gives similar results and is hence not reported.

Estimation error in expected returns is also accounted for with longer investment horizons. Results of the weekly and monthly rebalanced portfolios are consistent and reported in Table 6.9 and Table 6.10, respectively. It is obvious that the long memory volatility models produce portfolios that are more economically beneficial than the static and the short memory volatility models. Though the CGARCH-DCC and FIGARCH-DCC models generate the best results in some cases, their outperformance is not consistent and robust, especially that of the FIGARCH-DCC model. On the contrary, the performance of the LM-EWMA and LM-EWMA models is quite steady and reliable across different vectors of expected returns for the two datasets and across different investment horizons.

6.2.2 High Dimensional Systems: The International Stock and Bond and the DJIA Portfolios

As with the low dimensional systems, the whole sample of the high dimensional systems is divided into an estimation period and a forecast period. For the international stock and bond portfolio, the estimation period is from 1 Jan 1988 to 31 Dec 1993 (312 weekly observations) and the forecast period from 1 Jan 1994 to 31 Dec 2009 (835 observations). The estimation period of the DJIA portfolio ranges from 1 Mar 1990 to 29 Feb 1996 (1518 daily observations) and the forecast period from 1 Mar 1996 to 31 Dec 2009 (3483 observations). I, again, construct the optimal dynamic portfolios using the ex ante mean of the estimation period and rebalance the portfolios periodically based on changes in forecasts of the conditional covariance matrix.

Table 6.11 evaluates the out of sample performance of the weekly and monthly rebalanced international stock and bond portfolios. The FIGARCH-DCC model is excluded as it requires a prohibitively high upper lag cut-off. Similar results to the bivariate portfolios are identified. The dynamic strategies consistently outperform the static and the equally weighted strategies with all performance measures and rebalancing frequencies. The dynamic portfolios are generally more risky than the static portfolios, but they generate much higher returns, hence yielding very high Sharpe ratios. For example, the Sharpe ratios of the LM-EWMA portfolio are around 9 times as much as those of the static portfolio. Consequently, the annualised abnormal returns generated by the dynamic portfolios are relatively high, e.g., up to 3% with the LM-EWMA portfolio. The dynamic strategies also yield large performance fees due to their high realised returns. A week trader would be willing to pay 805 to 1242 *bps* to switch from the static to the dynamic LM-EWMA strategy. Moving from weekly to monthly rebalancing frequencies improves the Sharpe ratio, abnormal return and breakeven transaction cost. Since the weekly rebalanced dynamic portfolios are more risky, relatively to the static portfolio, than the monthly rebalanced dynamic portfolios, the passive investor with high risk tolerance $(\gamma = 1)$ is willing to pay more to switch at weekly rebalancing than at monthly rebalancing frequency. Conversely, the more risk-averse investor $(\gamma = 5)$ readily pays more when he rebalances monthly. Among the conditional volatility models, the LM-EWMA, LM-EWMA-DCC and EWMA models outperform significantly. While incorporating the long memory volatility dynamics in the DCC models significantly enhances portfolio performance, it is not so in the case of the non-DCC models. The trade-off is balanced between parsimony and correct specification; the difference between EWMA and LM-EWMA models is negligible in terms of Sharpe ratios and abnormal returns. Although it generally pays more to switch to the EWMA than to the LM-EWMA portfolios, it costs more in terms of transaction costs to implement the EWMA portfolio.

Performance of the DJIA portfolio is reported in Table 6.12. Consistent with the previous results, the investor is better off with the dynamic strategies than he is with the static and the equally-weighted strategies. When transaction costs are taken into consideration, however, the dynamic strategies are only attractive for lower rebalancing frequencies. Allowing for long memory volatility dynamics in forecasting the covariance matrix brings substantial gains, again with the exception of the FIGARCH-DCC model. The long memory LM-EWMA model generally produces superior portfolios to the short memory EWMA model, while the LM-EWMA-DCC and the CGARCH-DCC models systematically dominate the GARCH-DCC model. It is interesting to see the DCC structure performs remarkably well in this moderate correlation, high dimensional portfolio. In particular, the DCC models consistently outperform the simple cross product non-DCC models (EWMA and LM-EWMA models), which have so far been the best performing models. Also, owing to its parsimony, the LM-EWMA-DCC model consistently produces portfolios with the highest Sharpe ratios, abnormal returns and performance fees at all horizons among the DCC models. The LM-EWMA-DCC portfolio is also the most feasible portfolio in terms of transaction costs. Note that in Chapter 5, the DCC models produce forecasts of the covariance matrix that are more statistically accurate and economically useful than the non-DCC models in high dimensional systems. The results here confirm the benefits of the DCC models for high dimensional systems in the volatility timing framework. The findings also identify the good performance of the DCC models in high correlation systems. The greater flexibility that arises from separately estimating volatility and correlation is again beneficial in the high dimensional and/or high correlation systems. It would be of interest to examine this issue further in future research.

Estimation Error in Expected Returns

To account for estimation error in expected returns, I follow Fleming et al.'s (2001) recommendation to consider a range of expected returns that are generated via a bootstrap procedure. An artificial sample of 4000 observations is created by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. I then estimate the unconditional mean and covariance matrix of this artificial return series. Dynamic portfolios are constructed using the constant unconditional expected returns from the bootstrap and forecasts of the conditional covariance matrix. To ensure the static and the dynamic portfolios are based on the same ex ante information, the static benchmark portfolio is formed using the bootstrap constant expected returns and covariance matrix. I repeat this procedure with 1000 trials, studying the economic gains of volatility timing across a wide range of plausible vectors of expected returns.

Table 6.13 summarises the average results across the 1000 bootstrap vectors of expected returns for the international stock and bond portfolio. As with the previous results, it is obvious that the investor is willing to pay to switch from the static to the dynamic strategies, and from the short memory to the long memory volatility timing portfolios. The LM-EWMA model consistently produces the best portfolios in terms of all performance measures and across all investment horizons. Meanwhile, the LM-EWMA-DCC model is also the best among the DCC models. The LM-EWMA portfolio outperforms the static portfolio in terms of the Sharpe ratio in more than 70% of all bootstrap vectors of returns. The investor is also better off by at least 300 *bps* and up to 900 *bps*, on average, when switching to the LM-EWMA model. Evidence is similar for the DJIA portfolios, in which the LM-EWMA-DCC model yields the most superior results (Table 6.14). Note that the performance of the LM-EWMA model improves significantly when estimation error in expected returns is accounted for. Unlike the previous dismal results, the LM-EWMA portfolio is now the second best portfolio at both daily and weekly rebalancing frequencies. The other non-DCC short memory EWMA model, however, still produces poor portfolios which are even dominated by the static portfolios. The economic gains of the dynamic strategies are also evaluated for a quarterly investment horizon and the findings are similar.

The benefits of exploiting long memory vs. short memory volatility are now examined in greater detail. In particular, I evaluate the performance of the two best performing long memory and short memory volatility models in each dataset. Figure 6.1 shows the gains of employing the long memory LM-EWMA model relative to using the short memory EWMA model in the international stock and bond portfolio. The figure plots the realised Sharpe ratios for 1,000 trials of the bootstrap experiment. Each dot represents a separate trial, plotting the realised Sharpe ratios for both the short memory and the long memory volatility timing portfolios. The Sharpe ratio's distributions are clearly above the 45-degree line, suggesting the outperformance of the long memory volatility timing strategy in both rebalancing frequencies. An experiment with the DJIA portfolio produces similar results, where portfolios constructed with the long memory LM-EWMA-DCC model yield higher Sharpe ratios than portfolios constructed with the short memory GARCH-DCC model in nearly 84% of total trials (Figure 6.2).

Sensitivity to Risk Aversion Coefficient

In this section, I evaluate the performance of the dynamic strategies, controlling for different risk aversion coefficients λ. So far all reported results are based on $\lambda = 1$. For each value of the risk aversion coefficients λ, I again generate 1,000 bootstrap vectors of expected returns and use them, along with the conditional covariance matrix estimates, to construct the optimal portfolios. Table 6.15 and Table 6.16 evaluate the performance of the dynamic long memory and short memory volatility timing strategies against the static strategies for the two datasets. In particular, I evaluate the LM-EWMA model against the EWMA model in the international stock and bond portfolio, and the LM-EWMA-DCC model against the GARCH-DCC model in the DJIA portfolio. Not surprisingly, when the investor is more risk averse, he will choose portfolios with lower risk, accepting lower expected returns, and paying lower performance fees. The Sharpe ratios are approximately the same for all risk aversion coefficients, with the slight difference due to the bootstrap procedure. Again, the dynamic strategies – both long memory and short memory- consistently dominate the static strategies in both datasets with all rebalancing frequencies and risk aversion coefficients. In the international stock and bond portfolio, the long memory volatility timing portfolios yield higher Sharpe ratios than the static portfolios in over 70% of total trials. Though the long memory LM-EWMA portfolios also generally yield higher Sharpe ratios than the short memory EWMA portfolios, they underperform their short memory counterparts in terms of performance fees due to lower realised returns. In the DJIA portfolio, the long memory strategies, however, dominate the static and the short memory strategies by all performance measures with all risk aversion coefficients and across all investment horizons.

6.2.3 Sub-period Performance

Sub-period performance is of interest as the dynamic strategies can be evaluated through boom and bust periods. I study the sub-period performance of all four datasets, comparing the out of sample performance of the static strategy and the dynamic short memory and long memory strategies over the years. Following Han (2006), I calculate the yearly performance over the period from the beginning of the testing period (January 2, 1994) to the end of the target year. Some interesting observations emerge. First, our optimal portfolios, both static and dynamic, closely track the health of the economy. For example, the US portfolio returns decreased significantly after the dotcom crash in 2000 with a substantial increase in volatility, leading to a big drop in the Sharpe ratios. Then returns slowly increased, and so did the Sharpe ratios, but not back to the high level of the 1990s, before dropping again in 2008 when the recent recession started (Figure 6.3). Second, the dynamic strategies, especially the long memory volatility timing strategy, generally produce superior portfolios to the static strategy, especially in recession periods. The Sharpe ratios of both strategies evidently decreased when the market declined, yet the Sharpe ratio of the long memory volatility timing strategy dropped less markedly. The investor is also willing to pay more to switch to the dynamic strategies in recession periods than in normal periods. It may be inferred that the conditional volatility models better estimate high volatility in recession periods, helping the investor successfully control for the negative market changes by timing volatility in his investment decision. The results are consistent, even when accounting for estimation error in expected returns. To save space, only the average year-on-year performance of the static and the long memory LM-EWMA strategies (with bootstrap vectors of expected returns) for the international stock and bond portfolio is reported (Table 6.17).

6.2.4 An Additional Benchmark

An additional benchmark is employed to evaluate the performance of the long memory volatility timing portfolios. A naive dynamic strategy is constructed using the rolling window equally weighted covariance matrix estimator. The rolling window equally weighted covariance matrix accounts for time-varying volatility across different rolling windows; however it places equal weights on recent and distant observations in the same window. The dynamic strategy based on the rolling window covariance matrix is found to outperform the static strategy. However, the long memory volatility models still time volatility better in this experiment. Table 6.18 compares the performance of the rolling window and the long memory LM-EWMA portfolios. At all investment horizons, the long memory strategy generally yields higher Sharpe ratios than the rolling window strategy, especially in the low correlation Stock-Bond and the moderate correlation international stock and bond portfolios. For example, the investor is willing to pay an annualised performance fee τ_1 of up to 686 *bps* to switch from the rolling weekly rebalanced international stock and bond portfolio to the corresponding long memory portfolio. Though the domination of the long memory strategy is not as clearly marked in the high correlation portfolios, it is still beneficial to switch to the LM-EWMA model. Note that the gains from the long memory volatility timing strategy increase substantially in the higher correlation portfolios, especially in the high dimensional DJIA portfolio, when the LM-EWMA-DCC model is employed instead.

6.2.5 Long-Only Constraints

Short selling is generally not a common practice for most investors, and so in this section I evaluate the performance of the optimal portfolios under a long-only constraint. Besides, constraints are argued to be useful for controlling portfolio weights, hence reducing estimation error (see Frost and Savarino, 1988, Jagannathan and Ma, 2003). The findings are summarised in Table 6.19. To save space, I only report the results of the static portfolio, the short memory GARCH-DCC portfolio and the two long memory LM-EWMA and LM-EWMA-DCC portfolios. Compared to the previous results when no nonnegative weight constraint is applied, the performance of the two bivariate portfolios under the constraint is the same. The optimiser may mostly choose non-negative weights for the two bivariate portfolios even when no constraint is imposed. However, the results of the high dimensional portfolios are markedly different. Under the long-only constraint, portfolio performance, especially of the dynamic strategies, deteriorates significantly. For example, the Sharpe ratio of the weekly rebalanced LM-EWMA international and bond portfolio under the weight constraint is merely 0.3, as compared to that of 0.9 without the constraint. However, the long memory portfolios still dominate the static and the short memory portfolios in most cases. Again, the LM-EWMA model performs quite poorly, especially in the DJIA portfolio, where it is even dominated by the static unconditional alternative. The LM-EWMA-DCC model, on the contrary, produces consistently favourable results and so does the CGARCH-DCC model (results are not reported here). Although the use of a long-only constraint makes the dynamic strategies less attractive, it reduces the turnover of the dynamic portfolios, making them easier to be implemented in practice.

6.3 Conclusion

This chapter examines the economic value of allowing for long memory volatility in forecasting the conditional covariance matrix for dynamic asset allocation. Consistent with the literature, the results clearly demonstrate that investors are willing to pay to switch from the static unconditional strategy to the dynamic volatility timing alternatives. The findings also suggest that better volatility forecasts lead to better investment decisions. The long memory volatility timing portfolios consistently dominate the short memory volatility timing portfolios with all performance measures and across all investment horizons, with the exception of the FIGARCH-DCC portfolio. The high degree of parameterisation and computational burden may generate such high estimation error that it is detrimental to the performance of the FIGARCH-DCC model. The advantage of a parsimonious model is also proved by the consistent dominance of the LM-EWMA and LM-EWMA-DCC models among the long memory models. When transaction costs are considered, however, the dynamic volatility timing strategies are only attractive at lower rebalancing frequencies. The results apply to all four datasets and are robust to estimation error in expected returns, the choice of risk aversion coefficient, sub-period performance, benchmark strategies and the use of a long-only constraint.

The economic value of the conditional covariance matrix can be evaluated in some other directions. First, it would be of interest to extend the analysis in the context of time-varying expected returns. Here, expected returns are simply and unrealistically assumed constant. However, time-varying volatility affects returns and it is hence not justifiable to separate the movement of expected returns with those of volatility and correlation. Second, it may be useful to study the economic value of dynamic strategies in an intertemporal asset allocation framework. Dynamic strategies may behave differently in the presence of hedging demands. Third, the study is limited to evaluating the economic value of the long memory conditional covariance matrix from the perspective of a risk-averse investor. One may want to examine the implications of the long memory conditional covariance matrix in other practical problems, e.g., in risk management.

The poor performance of the non-DCC conditional covariance matrix estimators in the multivariate portfolios may point to another direction for future research. In Chapter 5, the non-DCC models have been found to produce poorer forecasts than the DCC models in high dimensional systems, and their underperformance in forecast power, in turn, leads to poorer volatility timing performance in asset allocation. It would be interesting to investigate this issue in greater detail.

Table 6.1. The Economic Values of Dynamic Strategies: Daily Rebalanced Bivariate Portfolios

The table compares the out-of-sample performance of the daily rebalanced bivariate portfolios. Panel A reports results for the Stock-Bond portfolio, and panel B for the S&P500-DJIA portfolio. 1/N is the equally weighted portfolio. The static portfolio is constructed using the constant mean and covariance matrix of the estimation period. For each dynamic strategy, the table reports the average annualised realised return (μ), the annualised realised volatility (σ), the Sharpe ratio (SR), the annualised abnormal return to the static portfolio (M2), the average annualised performance fee (in basis points) Δ_γ that an investor with a constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio.

	μ (%)	*σ (%)*	***SR***	***M2 (%)***	Δ_1	Δ_5
Panel A. Stock-Bond Portfolio						
1/N	3.127	10.314	0.303			
Static	0.267	0.869	0.308			
Volatility Timing Portfolios						
EWMA	0.684	1.229	0.557	0.216	41	40
GARCH-DCC	0.357	0.921	0.387	0.069	9	9
LM-EWMA	0.622	1.072	0.580	0.236	35	34
LM-EWMA-DCC	0.512	0.959	0.534	0.197	24	24
CGARCH-DCC	0.421	0.916	0.460	0.132	15	15
FIGARCH-DCC	0.101	2.642	0.038	-0.234	-20	-32
Panel B. S&P500-DJIA Portfolio						
1/N	2.395	19.180	0.125			
Static	0.078	0.516	0.152			
Volatility Timing Portfolios						
EWMA	0.155	0.530	0.292	0.072	8	8
GARCH-DCC	0.110	0.454	0.243	0.047	3	3
LM-EWMA	0.135	0.466	0.291	0.072	6	6
LM-EWMA-DCC	0.146	0.465	0.313	0.083	7	7
CGARCH-DCC	0.115	0.476	0.241	0.046	4	4
FIGARCH-DCC	-0.363	2.808	-0.129	-0.145	-48	-63

Table 6.2. The Economic Values of Dynamic Strategies: Weekly Rebalanced Bivariate Portfolios

The table compares the out-of-sample performance of the weekly rebalanced bivariate portfolios. Panel A reports results for the Stock-Bond portfolio, and panel B for the S&P500-DJIA portfolio. 1/N is the equally weighted portfolio. The static portfolio is constructed using the constant mean and covariance matrix of the estimation period. For each dynamic strategy, the table reports the average annualised realised return (μ), the annualised realised volatility (σ), the Sharpe ratio (SR), the annualised abnormal return to the static portfolio (M2), the average annualised performance fee (in basis points) Δ_γ that an investor with a constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio.

	μ (%)	σ (%)	*SR*	*M2 (%)*	Δ_1	Δ_5
Panel A. Stock-Bond Portfolio						
1/N	3.170	9.415	0.337			
Static	0.267	0.794	0.337			
Volatility Timing Portfolios						
EWMA	0.680	1.207	0.563	0.180	41	39
GARCH-DCC	0.396	0.850	0.466	0.102	13	13
LM-EWMA	0.579	1.016	0.570	0.185	31	30
LM-EWMA-DCC	0.515	0.908	0.567	0.183	25	24
CGARCH-DCC	0.486	0.846	0.574	0.188	22	22
FIGARCH-DCC	0.357	1.586	0.225	-0.089	8	4
Panel B. S&P500-DJIA Portfolio						
1/N	2.463	17.637	0.140			
Static	0.078	0.471	0.166			
Volatility Timing Portfolios						
EWMA	0.162	0.543	0.297	0.062	8	8
GARCH-DCC	0.122	0.442	0.277	0.052	4	4
LM-EWMA	0.145	0.465	0.311	0.068	7	7
LM-EWMA-DCC	0.156	0.464	0.336	0.080	8	8
CGARCH-DCC	0.121	0.464	0.261	0.045	4	4
FIGARCH-DCC	0.096	1.460	0.065	-0.048	1	-3

Table 6.3. The Economic Values of Dynamic Strategies: Monthly Rebalanced Bivariate Portfolios

The table compares the out-of-sample performance of the monthly rebalanced bivariate portfolios. Panel A reports results for the Stock-Bond portfolio, and panel B for the S&P500-DJIA portfolio. 1/N is the equally weighted portfolio. The static portfolio is constructed using the constant mean and covariance matrix of the estimation period. For each dynamic strategy, the table reports the average annualised realised return (μ), the annualised realised volatility (σ), the Sharpe ratio (SR), the annualised abnormal return to the static portfolio (M2), the average annualised performance fee (in basis points) Δ_γ that an investor with a constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio.

	μ (%)	σ (%)	SR	M2 (%)	Δ_1	Δ_5
Panel A. Stock-Bond Portfolio						
1/N	3.219	8.900	0.362			
Static	0.273	0.742	0.367			
Volatility Timing Portfolios						
EWMA	0.726	1.091	0.666	0.221	45	44
GARCH-DCC	0.449	0.803	0.559	0.142	18	17
LM-EWMA	0.659	0.987	0.667	0.222	38	38
LM-EWMA-DCC	0.588	0.862	0.683	0.234	31	31
CGARCH-DCC	0.519	0.814	0.638	0.201	25	24
FIGARCH-DCC	0.827	1.244	0.665	0.221	55	53
Panel B. S&P500-DJIA Portfolio						
1/N	2.398	16.623	0.144			
Static	0.078	0.457	0.170			
Volatility Timing Portfolios						
EWMA	0.192	0.518	0.370	0.092	11	11
GARCH-DCC	0.150	0.452	0.331	0.074	7	7
LM-EWMA	0.165	0.465	0.355	0.085	9	9
LM-EWMA-DCC	0.172	0.468	0.367	0.090	9	9
CGARCH-DCC	0.161	0.468	0.343	0.079	8	8
FIGARCH-DCC	0.165	0.947	0.174	0.002	8	7

Table 6.4. Performance Fees to Switch from the Short Memory GARCH Volatility Timing Strategy to the Long Memory Volatility Timing Strategy

The table measures the average annualised performance fees (in basis points) Δ_γ that an investor with a constant relative risk coefficient of γ is willing to pay to switch from the portfolio constructed with the short memory GARCH model to the portfolios constructed with the long memory volatility models in the first column.

	Stock-Bond Portfolio		***S&P500-DJIA Portfolio***	
	Δ_1	Δ_5	Δ_1	Δ_5
Panel A. Daily Rebalancing				
LM-EWMA	26	26	3	2
LM-EWMA-DCC	16	15	4	4
CGARCH-DCC	6	6	0	0
FIGARCH-DCC	-29	-41	-51	-67
Panel B. Weekly Rebalancing				
LM-EWMA	18	18	2	2
LM-EWMA-DCC	12	12	3	3
CGARCH-DCC	9	9	0	0
FIGARCH-DCC	-5	-8	-4	-8
Panel C. Monthly Rebalancing				
LM-EWMA	21	20	2	2
LM-EWMA-DCC	14	14	2	2
CGARCH-DCC	7	7	1	1
FIGARCH-DCC	37	36	1	0

Table 6.5. Performance Fees to Switch from the Short Memory EWMA Volatility Timing Strategy to the Long Memory Volatility Timing Strategy

The table measures the average annualised performance fees (in basis points) Δ_γ that an investor with a constant relative risk coefficient of γ is willing to pay to switch from the portfolio constructed with the short memory EWMA model to the portfolios constructed with the long memory volatility models in the first column.

	Stock-Bond Portfolio		***S&P500-DJIA Portfolio***	
	Δ_1	Δ_5	Δ_1	Δ_5
Panel A. Daily Rebalancing				
LM-EWMA	-6	-5	-2	-2
LM-EWMA-DCC	-17	-16	-1	-1
CGARCH-DCC	-26	-25	-4	-4
FIGARCH-DCC	-61	-72	-56	-71
Panel B. Weekly Rebalancing				
LM-EWMA	-10	-9	-2	-1
LM-EWMA-DCC	-16	-15	-1	0
CGARCH-DCC	-19	-18	-4	-4
FIGARCH-DCC	-33	-35	-8	-11
Panel C. Monthly Rebalancing				
LM-EWMA	-7	-6	-3	-3
LM-EWMA-DCC	-14	-13	-2	-2
CGARCH-DCC	-20	-19	-3	-3
FIGARCH-DCC	10	9	-3	-4

Table 6.6. Breakeven Transaction Costs of the Bivariate Portfolios

The table reports the breakeven transaction costs τ_{γ} (in basis points) for daily, weekly and monthly rebalanced portfolios. Panel A reports results for the Stock-Bond portfolio, and panel B for the S&P500-DJIA portfolio. For each dynamic strategy, if an investor with a relative risk coefficient γ has a transaction cost lower than the breakeven transaction cost τ_{γ}, he will be better off with the dynamic strategy; otherwise he should follow the static strategy. The breakeven transaction costs are only estimated when the performance fees are positive.

	Daily rebalancing		***Weekly rebalancing***		***Monthly rebalancing***	
	τ_1	τ_5	τ_1	τ_5	τ_1	τ_5
Panel A. Stock-Bond Portfolio						
EWMA	11	10	20	19	79	77
GARCH-DCC	2	2	11	11	40	40
LM-EWMA	9	9	19	19	75	73
LM-EWMA DCC	5	5	17	17	70	70
CGARCH-DCC	3	3	18	17	57	57
FIGARCH-DCC	–	–	1	1	40	38
Panel B. S&P500-DJIA portfolio						
EWMA	3	3	6	6	34	34
GARCH-DCC	1	1	5	5	19	19
LM-EWMA	2	2	6	6	28	28
LM-EWMA DCC	2	2	7	7	27	27
CGARCH-DCC	1	1	3	3	18	18
FIGARCH-DCC	–	–	0	–	4	3

Table 6.7. Estimation Error in Expected Returns: The Sharpe Ratios of the Daily Rebalanced Bivariate Portfolios

The table reports the out-of-sample Sharpe ratios of the daily rebalanced bivariate portfolios constructed with all possible pairs of expected returns. Panel A reports the Sharpe ratios for the Stock-Bond portfolio, and panel B for the S&P500-DJIA portfolio. Each row in the table reports the results for the pair of expected returns of the corresponding first two columns. The weighted returns are estimated using the Bayesian prior probabilities.

Panel A. Stock-Bond Portfolio

μ_{Stock}	μ_{Bond}	*EWMA*	*GARCH DCC*	*LM EWMA*	*LM EWMA DCC*	*CGARCH DCC*	*FIGARCH DCC*	*Static*
0.00	1.00	0.227	-0.057	0.166	0.093	-0.022	0.243	0.016
0.16	0.99	0.261	-0.026	0.204	0.127	0.012	0.239	0.035
0.31	0.95	0.296	0.008	0.243	0.163	0.049	0.229	0.058
0.45	0.89	0.333	0.046	0.285	0.204	0.090	0.213	0.085
0.59	0.81	0.373	0.090	0.331	0.250	0.139	0.190	0.121
0.71	0.71	0.417	0.143	0.383	0.304	0.197	0.161	0.168
0.81	0.59	0.465	0.208	0.443	0.370	0.269	0.126	0.230
0.89	0.45	0.515	0.286	0.508	0.446	0.355	0.088	0.294
0.95	0.31	0.553	0.371	0.570	0.522	0.444	0.047	0.313
0.99	0.16	0.551	0.437	0.599	0.556	0.500	0.006	0.271
1.00	0.00	0.479	0.451	0.556	0.507	0.490	-0.034	0.217
Weighted		*0.549*	*0.360*	*0.563*	*0.513*	*0.433*	*0.052*	*0.314*

Panel B. S&P500-DJIA Portfolio

$\mu_{S\&P500}$	μ_{DJIA}	*EWMA*	*GARCH DCC*	*LM EWMA*	*LM EWMA DCC*	*CGARCH DCC*	*FIGARCH DCC*	*Static*
0.00	1.00	-0.002	0.076	-0.011	0.056	0.030	0.243	0.164
0.16	0.99	0.017	0.089	0.008	0.074	0.043	0.239	0.172
0.31	0.95	0.045	0.109	0.036	0.101	0.063	0.229	0.184
0.45	0.89	0.093	0.140	0.084	0.146	0.098	0.213	0.204
0.59	0.81	0.186	0.196	0.180	0.230	0.165	0.190	0.232
0.71	0.71	0.318	0.244	0.317	0.323	0.255	0.161	0.092
0.81	0.59	0.272	0.156	0.273	0.228	0.188	0.126	-0.043
0.89	0.45	0.195	0.080	0.198	0.139	0.118	0.088	-0.085
0.95	0.31	0.151	0.041	0.156	0.092	0.082	0.047	-0.104
0.99	0.16	0.124	0.019	0.130	0.065	0.061	0.006	-0.115
1.00	0.00	0.105	0.004	0.112	0.046	0.047	-0.034	-0.123
Weighted		*0.316*	*0.245*	*0.314*	*0.323*	*0.254*	*0.052*	*0.099*

Table 6.8. Estimation Error in Expected Returns: Relative Performance Fees of the Daily Rebalanced Bivariate Portfolios

The table reports the relative performance fees that a day trader with a constant relatively risk aversion of 1 is willing to pay to switch from the static strategy to the dynamic trading strategies. Panel A reports the fees for the Stock-Bond portfolio, and panel B for the S&P500-DJIA portfolio. Each row in the table reports the fees for the pair of expected returns of the corresponding first two columns. For each vector of expected return, the highest positive performance fee is normalised to 100. A number of 95 indicates that the performance fee is equal to 95% of the highest performance fee. A "–" sign implies a negative performance fee and the investor will stay with the static strategy. The weighted returns are estimated using the Bayesian prior probabilities.

Panel A. Stock-Bond Portfolio

μ_{Stock}	μ_{Bond}	*EWMA*	*GARCH DCC*	*LM EWMA*	*LM-EWMA DCC*	*CGARCH DCC*	*FIGARCH DCC*
0.00	1.00	100	13	93	55	19	91
0.16	0.99	100	4	92	52	14	68
0.31	0.95	100	–	89	50	8	36
0.45	0.89	100	–	87	47	5	–
0.59	0.81	100	–	85	46	3	–
0.71	0.71	100	–	84	48	5	–
0.81	0.59	100	–	84	51	12	–
0.89	0.45	100	2	87	57	24	–
0.95	0.31	100	20	92	66	40	–
0.99	0.16	100	43	100	78	60	–
1.00	0.00	89	65	100	83	78	–
Weighted		*100*	*18*	*84*	*58*	*34*	–

Panel B. S&P500-DJIA Portfolio

$\mu_{S\&P500}$	μ_{DJIA}	*EWMA*	*GARCH DCC*	*LM EWMA*	*LM-EWMA DCC*	*CGARCH DCC*	*FIGARCH DCC*
0.00	1.00	25	100	55	79	49	–
0.16	0.99	–	100	36	74	29	–
0.31	0.95	–	100	–	61	–	–
0.45	0.89	–	–	–	–	–	–
0.59	0.81	–	–	–	–	–	–
0.71	0.71	100	64	92	94	70	–
0.81	0.59	100	79	99	91	84	–
0.89	0.45	97	86	100	91	87	–
0.95	0.31	95	89	100	92	88	–
0.99	0.16	93	91	100	92	88	–
1.00	0.00	92	93	100	92	88	–
Weighted		*100*	*51*	*83*	*86*	*59*	–

Table 6.9. Estimation Error in Expected Returns: The Sharpe Ratios of the Weekly Rebalanced Bivariate Portfolios

The table reports the out-of-sample Sharpe ratios of the weekly rebalanced bivariate portfolios with all possible pairs of expected returns. Panel A reports the Sharpe ratios for the Stock-Bond portfolio, and panel B for the S&P500-DJIA portfolio. Each row in the table reports the results for the pair of expected returns of the corresponding first two columns. The weighted returns are estimated using the Bayesian prior probabilities.

Panel A. Stock-Bond Portfolio

μ_{Stock}	μ_{Bond}	*EWMA*	*GARCH DCC*	*LM EWMA*	*LM EWMA DCC*	*CGARCH DCC*	*FIGARCH DCC*	*Static*
0.00	1.00	0.139	-0.038	0.060	0.078	0.000	0.081	0.017
0.16	0.99	0.173	-0.004	0.099	0.113	0.039	0.101	0.036
0.31	0.95	0.209	0.034	0.140	0.149	0.082	0.121	0.059
0.45	0.89	0.248	0.076	0.184	0.190	0.129	0.142	0.087
0.59	0.81	0.291	0.124	0.234	0.237	0.184	0.163	0.121
0.71	0.71	0.341	0.182	0.292	0.293	0.251	0.183	0.168
0.81	0.59	0.402	0.254	0.363	0.363	0.334	0.203	0.230
0.89	0.45	0.473	0.344	0.450	0.450	0.438	0.218	0.301
0.95	0.31	0.550	0.446	0.550	0.549	0.553	0.225	0.338
0.99	0.16	0.599	0.530	0.633	0.615	0.634	0.220	0.305
1.00	0.00	0.564	0.544	0.637	0.579	0.621	0.200	0.244
Weighted		*0.539*	*0.432*	*0.536*	*0.536*	*0.537*	*0.225*	*0.338*

Panel B. S&P500-DJIA Portfolio

$\mu_{S\&P500}$	μ_{DJIA}	*EWMA*	*GARCH DCC*	*LM EWMA*	*LM EWMA DCC*	*CGARCH DCC*	*FIGARCH DCC*	*Static*
0.00	1.00	-0.001	0.068	0.000	0.044	0.007	0.149	0.158
0.16	0.99	0.017	0.083	0.019	0.063	0.025	0.153	0.165
0.31	0.95	0.044	0.104	0.046	0.090	0.051	0.158	0.175
0.45	0.89	0.089	0.139	0.094	0.136	0.098	0.159	0.193
0.59	0.81	0.180	0.206	0.189	0.227	0.197	0.132	0.223
0.71	0.71	0.333	0.290	0.346	0.361	0.390	0.038	0.103
0.81	0.59	0.292	0.196	0.293	0.266	0.304	-0.039	-0.045
0.89	0.45	0.200	0.105	0.201	0.162	0.194	-0.076	-0.087
0.95	0.31	0.151	0.061	0.152	0.110	0.143	-0.094	-0.105
0.99	0.16	0.123	0.035	0.123	0.080	0.113	-0.105	-0.115
1.00	0.00	0.104	0.018	0.104	0.061	0.094	-0.113	-0.122
Weighted		*0.330*	*0.289*	*0.342*	*0.359*	*0.385*	*0.042*	*0.111*

Table 6.10. Estimation Error in Expected Returns: The Sharpe Ratios of the Monthly Rebalanced Bivariate Portfolios

The table reports the out-of-sample Sharpe ratios of the monthly rebalanced bivariate portfolios with all possible pairs of expected returns. Panel A reports the Sharpe ratios for the Stock-Bond portfolio, and panel B for the S&P500-DJIA portfolio. Each row in the table reports the results for the pair of expected returns of the corresponding first two columns. The weighted returns are estimated using the Bayesian priors probability.

Panel A. Stock-Bond Portfolio

μ_{Stock}	μ_{Bond}	*EWMA*	*GARCH DCC*	*LM EWMA*	*LM EWMA DCC*	*CGARCH DCC*	*FIGARCH DCC*	*Static*
0.00	1.00	0.120	0.001	0.108	0.139	0.095	0.006	0.044
0.16	0.99	0.159	0.039	0.148	0.174	0.128	0.066	0.065
0.31	0.95	0.199	0.080	0.191	0.211	0.163	0.129	0.089
0.45	0.89	0.242	0.125	0.236	0.253	0.203	0.197	0.118
0.59	0.81	0.292	0.178	0.288	0.301	0.250	0.273	0.154
0.71	0.71	0.351	0.241	0.350	0.361	0.308	0.359	0.203
0.81	0.59	0.424	0.320	0.426	0.437	0.384	0.456	0.268
0.89	0.45	0.519	0.420	0.523	0.537	0.487	0.557	0.341
0.95	0.31	0.639	0.536	0.642	0.659	0.613	0.650	0.372
0.99	0.16	0.760	0.628	0.754	0.742	0.698	0.705	0.321
1.00	0.00	0.773	0.626	0.763	0.663	0.620	0.699	0.245
Weighted		*0.622*	*0.520*	*0.624*	*0.642*	*0.596*	*0.639*	*0.373*

Panel B. S&P500-DJIA Portfolio

$\mu_{S\&P500}$	μ_{DJIA}	*EWMA*	*GARCH DCC*	*LM EWMA*	*LM EWMA DCC*	*CGARCH DCC*	*FIGARCH DCC*	*Static*
0.00	1.00	0.086	0.110	0.056	0.062	-0.018	-0.120	0.173
0.16	0.99	0.104	0.125	0.074	0.079	-0.003	-0.107	0.179
0.31	0.95	0.129	0.148	0.100	0.105	0.020	-0.085	0.189
0.45	0.89	0.172	0.185	0.144	0.150	0.059	-0.039	0.205
0.59	0.81	0.257	0.255	0.233	0.241	0.143	0.072	0.229
0.71	0.71	0.403	0.343	0.391	0.402	0.299	0.191	0.103
0.81	0.59	0.272	0.208	0.282	0.273	0.259	0.200	-0.057
0.89	0.45	0.135	0.090	0.157	0.148	0.181	0.189	-0.102
0.95	0.31	0.074	0.035	0.100	0.092	0.141	0.180	-0.120
0.99	0.16	0.041	0.005	0.068	0.061	0.118	0.173	-0.130
1.00	0.00	0.020	-0.014	0.048	0.041	0.102	0.167	-0.137
Weighted		*0.400*	*0.343*	*0.388*	*0.399*	*0.295*	*0.190*	*0.112*

Table 6.11. Portfolio Performance of the International Stock and Bond Portfolio

The table compares the out-of-sample performance of the optimal international stock and bond portfolio. Panels A and B report the results of the weekly and monthly rebalanced portfolios, respectively. The static portfolio is constructed using the constant mean and covariance matrix of the estimation period. For each dynamic strategy, the table reports the average annualised realised return (μ), the annualised realised volatility (σ), the Sharpe ratio (SR), the annualised abnormal return to the static portfolio (M2), the annualised performance fee (in basis points) Δ_γ that an investor with a constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio, and the breakeven transaction cost τ_γ (in basis points) that he will be better off with the dynamic strategy.

	μ (%)	*σ (%)*	***SR***	***M2 (%)***	Δ_1	Δ_5	τ_1	τ_5
Panel A. Weekly rebalancing								
1/N	3.911	12.996	-0.007					
Static	4.393	3.836	0.103					
Volatility timing strategies								
EWMA	48.263	46.839	0.945	3.232	3297	-1230	3	–
GARCH-DCC	12.446	20.077	0.421	1.221	611	-173	1	–
LM-EWMA	17.897	15.220	0.913	3.109	1242	805	7	4
LM-EWMA-DCC	13.382	17.319	0.542	1.685	756	181	1	0
CGARCH-DCC	12.067	16.521	0.488	1.480	638	118	2	0
Panel B. Monthly rebalancing								
1/N	4.071	14.064	0.005					
Static	4.354	3.874	0.092					
Volatility timing strategies								
EWMA	20.060	14.468	1.110	3.947	1474	1074	11	8
GARCH-DCC	9.954	12.512	0.476	1.489	489	199	3	1
LM-EWMA	15.938	11.537	1.035	3.655	1099	858	12	9
LM-EWMA-DCC	11.925	11.093	0.714	2.414	703	482	4	3
CGARCH-DCC	9.660	12.662	0.447	1.378	458	160	5	2

Table 6.12. Portfolio Performance of the DJIA Portfolio

The table compares the out-of-sample performance of the optimal DJIA portfolio. Panels A, B and C report the results of the daily, weekly, and monthly rebalanced portfolios, respectively. The static portfolio is constructed using the constant mean and covariance matrix of the estimation period. For each dynamic strategy, the table reports the average annualised realised return (μ), the annualised realised volatility (σ), the Sharpe ratio (SR), the annualised abnormal return to the static portfolio (M2), the annualised performance fee (in basis points) Δ_γ that an investor with a constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio, and the breakeven transaction cost τ_γ (in basis points) that he will be better off with the dynamic strategy.

	μ (%)	σ (%)	*SR*	*M2 (%)*	Δ_1	Δ_5	τ_1	τ_5
Panel A. Daily rebalancing								
1/N	5.428	21.191	0.067					
Static	4.166	1.796	0.092					
Volatility timing strategies								
EWMA	4.965	14.690	0.066	-0.048	-26	-452	0	–
GARCH-DCC	4.620	3.294	0.188	0.172	42	26	3	2
LM-EWMA	4.981	5.807	0.169	0.137	66	5	2	0
LM-EWMA-DCC	5.014	3.352	0.303	0.378	81	65	6	4
CGARCH-DCC	4.488	3.536	0.138	0.082	28	9	1	0
FIGARCH-DCC	3.363	8.823	-0.072	-0.296	-118	-267	–	–
Panel B. Weekly rebalancing								
1/N	5.509	19.265	0.078					
Static	4.170	1.660	0.102					
Volatility timing strategies								
EWMA	5.856	15.393	0.121	0.030	52	-420	1	–
GARCH-DCC	4.587	3.394	0.173	0.117	37	20	7	4
LM-EWMA	4.715	5.404	0.132	0.050	41	-12	2	–
LM-EWMA-DCC	4.688	3.470	0.198	0.159	47	28	8	5
CGARCH-DCC	4.560	3.536	0.158	0.093	34	15	7	4
FIGARCH-DCC	-1.712	18.124	-0.315	-0.693	-751	-1408	–	–
Panel C. Monthly rebalancing								
1/N	5.427	19.835	0.072					
Static	4.182	1.649	0.111					
Volatility timing strategies								
EWMA	4.518	6.277	0.082	-0.046	15	-59	2	–
GARCH-DCC	4.879	3.002	0.293	0.301	67	54	39	32
LM-EWMA	4.481	4.565	0.105	-0.009	21	-16	4	–
LM-EWMA-DCC	5.183	3.212	0.368	0.425	96	81	56	47
CGARCH-DCC	4.978	3.271	0.299	0.311	76	59	43	34
FIGARCH-DCC	3.227	6.275	-0.123	-0.386	-114	-188	–	–

Table 6.13. Average Portfolio Performance of the International Stock and Bond Portfolio with Bootstrap Experiments

The table reports the average out-of-sample performance of the international stock and bond portfolio across a wide range of bootstrap-generated expected returns. An artificial sample of 4000 observations is generated by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. The procedure is repeated with 1,000 trials. Panels A and B report the results of the weekly and monthly rebalanced portfolios, respectively. The static portfolios are constructed using the bootstrap expected returns and covariance matrices. For each dynamic strategy, the table reports the average annualised realised return (μ), the average annualised realised volatility (σ), the average Sharpe ratio (SR), the *p*-value (proportion) that the dynamic strategy outperforms the static alternative in terms of the Sharpe ratio, the average annualised abnormal return to the static portfolio (M2), the average annualised performance fee (in basis points) Δ_γ that an investor a constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio, and the average breakeven transaction cost τ_γ (in basis points) that he will be better off with the dynamic strategy.

	μ (%)	σ (%)	SR	p-value	M2 (%)	Δ_1	Δ_5	τ_1	τ_5
Panel A. Weekly rebalancing									
Static	5.557	3.408	0.458						
Volatility timing strategies									
EWMA	17.116	18.636	0.704	0.680	0.852	972	226	4	1
GARCH DCC	11.124	18.407	0.405	0.490	-0.176	374	-369	1	–
LM-EWMA	15.629	15.399	0.744	0.713	0.988	883	382	5	2
LM-EWMA DCC	11.715	16.207	0.489	0.562	0.113	476	-91	1	–
CGARCH DCC	9.995	14.588	0.426	0.508	-0.102	331	-125	1	–
Panel B. Monthly rebalancing									
Static	5.507	3.557	0.426						
Volatility timing strategies									
EWMA	16.987	17.064	0.764	0.712	1.225	993	343	7	3
GARCH-DCC	10.063	14.954	0.434	0.557	0.039	334	-178	2	–
LM-EWMA	14.606	12.713	0.832	0.740	1.463	827	486	8	5
LM-EWMA DCC	11.720	16.479	0.489	0.606	0.237	465	-200	2	–
CGARCH DCC	9.510	12.156	0.505	0.574	0.291	322	-2	3	–

Table 6.14. Average Portfolio Performance of the DJIA Portfolio with Bootstrap Experiments

The table reports the average out-of-sample performance of the international stock and bond portfolio across a wide range of bootstrap-generated expected returns. An artificial sample of 4000 observations is generated by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. The procedure is repeated with 1,000 trials. Panels A, B and C report results for the daily, weekly and monthly rebalanced portfolios, respectively. The static portfolios are constructed using the bootstrap expected returns and covariance matrices. For each dynamic strategy, the table reports the average annualised realised return (μ), the average annualised realised volatility (σ), the average Sharpe ratio (SR), the p-value (proportion) that the dynamic strategy outperforms the static alternative in terms of the Sharpe ratio, the average annualised abnormal return to the static portfolio (M2), the average annualised performance fee (in basis points) Δ_γ that an investor with a constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio, and the average breakeven transaction cost τ_γ (in basis points) that he will be better off with the dynamic strategy.

	μ (%)	σ (%)	*SR*	*p-value*	*M2 (%)*	Δ_1	Δ_5	τ_1	τ_5
Panel A. Daily rebalancing									
Static	4.620	3.197	0.201						
Volatility timing strategies									
EWMA	6.268	12.539	0.180	0.460	-0.045	90	-211	1	–
GARCH-DCC	4.752	2.824	0.264	0.706	0.217	15	20	1	1
LM-EWMA	5.489	5.080	0.290	0.726	0.307	79	48	2	1
LM-EWMA-DCC	4.981	2.900	0.336	0.824	0.453	37	42	3	3
CGARCH-DCC	4.802	2.948	0.269	0.697	0.230	19	23	1	1
FIGARCH-DCC	4.295	8.474	0.051	0.184	-0.497	-66	-199	–	–
Panel B. Weekly rebalancing									
Static	4.634	2.856	0.230						
Volatility timing strategies									
EWMA	5.248	5.748	0.215	0.466	-0.024	49	-2	3	–
GARCH-DCC	4.702	2.766	0.252	0.576	0.080	7	9	1	2
LM-EWMA	5.270	4.739	0.265	0.592	0.118	56	28	3	2
LM-EWMA-DCC	4.772	2.969	0.260	0.590	0.111	14	13	2	2
CGARCH-DCC	4.788	2.846	0.276	0.622	0.143	16	16	3	3
FIGARCH-DCC	2.386	10.994	-0.129	0.023	-1.025	-288	-543	–	–
Panel C. Monthly rebalancing									
Static	4.629	2.815	0.234						
Volatility timing strategies									
EWMA	5.145	5.457	0.211	0.442	-0.031	41	-5	5	–
GARCH-DCC	4.750	2.453	0.306	0.693	0.229	13	18	8	11
LM-EWMA	5.140	3.772	0.302	0.664	0.219	48	35	10	7
LM-EWMA-DCC	4.967	2.516	0.387	0.830	0.467	35	39	21	24
CGARCH-DCC	4.851	2.639	0.322	0.736	0.276	23	25	14	15
FIGARCH-DCC	4.324	5.225	0.067	0.107	-0.459	-40	-81	–	–

Table 6.15. Comparison of the Static and the Dynamic Volatility Timing Strategies Using Different Risk Aversion Coefficients: International Stock and Bond Portfolio

The table compares the average out-of-sample performance of the static and dynamic strategies with different risk aversion coefficients λ. A bootstrap procedure is applied to account for estimation error in expected returns. An artificial sample of 4000 observations is generated by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. The procedure is repeated with 1,000 trials. Panels A and B report the results of the weekly and monthly rebalanced portfolios, respectively. The static portfolios are constructed using the bootstrap expected returns and covariance matrices. The short memory portfolios are constructed with the EWMA model, while the long memory model portfolios are constructed with the LM-EWMA model. For each dynamic strategy, the table reports the average annualised realised return (μ), the average annualised realised volatility (σ), the average Sharpe ratio (SR), the p-value (proportion) that the dynamic strategy outperforms the static alternative in terms of the Sharpe ratio, and the average annualised performance fee (in basis points) Δ_γ that an investor with a constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio.

	Static			***Short memory***			***Short memory vs. Static***			***Long memory***			***Long memory vs. Static***		
λ	***μ (%)***	***σ (%)***	***SR***	***μ (%)***	***σ (%)***	***SR***	***p-value***	Δ_1	Δ_5	***μ (%)***	***σ (%)***	***SR***	***p-value***	Δ_1	Δ_5
Panel A. Weekly rebalancing															
1	5.557	3.408	0.458	17.116	18.636	0.704	0.680	972	226	15.629	15.399	0.744	0.713	883	382
2	4.783	1.714	0.458	10.825	9.336	0.728	0.697	558	369	10.033	7.715	0.766	0.729	493	367
3	4.522	1.144	0.458	8.501	6.190	0.724	0.688	378	296	7.998	5.156	0.760	0.730	334	277
4	4.391	0.853	0.460	7.306	4.709	0.707	0.698	280	232	6.909	3.887	0.742	0.725	244	212
5	4.313	0.685	0.458	6.718	3.759	0.731	0.703	233	203	6.406	3.110	0.770	0.741	204	184
Panel B. Monthly rebalancing															
1	5.507	3.557	0.426	16.987	17.064	0.764	0.712	993	343	14.606	12.713	0.832	0.740	827	486
2	4.758	1.790	0.425	10.760	8.514	0.797	0.730	561	402	9.505	6.321	0.865	0.758	454	370
3	4.506	1.194	0.425	8.458	5.637	0.795	0.723	378	309	7.644	4.219	0.860	0.755	305	268
4	4.379	0.890	0.428	7.269	4.305	0.770	0.727	279	239	6.650	3.206	0.831	0.750	222	200
5	4.303	0.714	0.425	6.692	3.451	0.791	0.735	233	207	6.196	2.571	0.858	0.764	186	172

Table 6.16. Comparison of the Static and the Dynamic Volatility Timing Strategies Using Different Risk Aversion Coefficients: DJIA Portfolio

The table compares the average out-of-sample performance of the static and dynamic strategies using different risk aversion coefficients λ. A bootstrap procedure is applied to control for estimation error in expected returns. The static portfolios are constructed using the bootstrap expected returns and covariance matrices. The short memory portfolios are constructed with the GARCH-DCC model, while the long memory portfolios are constructed with the LM-EWMA-DCC model. For each dynamic strategy, the table reports the average annualised realised return (μ), the annualised realised volatility (σ), the Sharpe ratio (SR), the *p*-value (proportion) that the dynamic strategy outperforms the static alternative in terms of the Sharpe ratio, and the annualised performance fee (in basis points) Δ_γ that an investor with a relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio.

	Static			*Short memory*			*Short memory vs. Static*			*Long memory*			*Long memory vs. Static*		
λ	*μ (%)*	*σ (%)*	*SR*	*μ (%)*	*σ (%)*	*SR*	*p-value*	Δ_1	Δ_5	*μ (%)*	*σ (%)*	*SR*	*p-value*	Δ_1	Δ_5
Panel A. Daily rebalancing															
1	4.620	3.197	0.201	4.752	2.824	0.264	0.706	15	20	4.981	2.900	0.336	0.824	37	42
2	4.312	1.613	0.199	4.375	1.417	0.263	0.694	7	8	4.492	1.454	0.335	0.830	18	19
3	4.204	1.060	0.200	4.250	0.937	0.265	0.719	5	5	4.328	0.962	0.337	0.848	12	13
4	4.153	0.794	0.200	4.185	0.705	0.260	0.688	3	4	4.244	0.724	0.334	0.827	9	9
5	4.123	0.636	0.200	4.149	0.560	0.264	0.704	3	3	0.704	2.6	2.9	0.816	7	7
Panel B. Weekly rebalancing															
1	4.634	2.856	0.230	4.702	2.766	0.252	0.576	7	9	4.772	2.969	0.260	0.590	14	13
2	4.320	1.442	0.228	4.350	1.386	0.250	0.579	3	4	4.380	1.494	0.255	0.573	6	6
3	4.209	0.947	0.229	4.233	0.916	0.253	0.592	2	3	4.257	0.985	0.261	0.596	5	5
4	4.157	0.709	0.230	4.174	0.690	0.250	0.560	2	2	4.191	0.743	0.258	0.553	3	3
5	4.126	0.568	0.230	4.139	0.549	0.251	0.580	1	1	4.151	0.591	0.256	0.584	2	2
Panel C. Monthly rebalancing															
1	4.629	2.815	0.234	4.750	2.453	0.306	0.693	13	18	4.967	2.516	0.387	0.830	35	39
2	4.317	1.421	0.232	4.377	1.232	0.305	0.690	6	7	4.487	1.262	0.387	0.832	17	18
3	4.207	0.933	0.233	4.252	0.812	0.310	0.704	5	5	4.325	0.832	0.392	0.852	12	12
4	4.155	0.699	0.233	4.191	0.615	0.308	0.689	4	4	4.245	0.629	0.390	0.835	9	9
5	4.125	0.559	0.233	4.152	0.489	0.310	0.706	3	3	4.194	0.500	0.389	0.826	7	7

Table 6.17. Yearly Performance of the International Stock and Bond Portfolio

The table reports the average yearly performance of the international stock and bond portfolio. A bootstrap procedure is applied to control for estimation error in expected returns. The static portfolios are constructed using the bootstrap expected returns and covariance matrices, while the long memory volatility timing portfolios are constructed based on the bootstrap expected returns and forecasts of the conditional covariance matrix from the LM-EWMA model. The yearly performance is calculated over the period from the beginning of the testing period (January 2, 1994) to the end of the target year. The table reports the average annualised realised returns (μ), the annualised realised volatilities (σ), the Sharpe ratios (SR), the *p*-values (proportion) that the dynamic strategies outperform the static strategies in terms of the Sharpe ratio, and the average annualised performance fees Δ_γ (in basis points) that an investor with a constant relative risk coefficient of γ is willing to pay to switch from the static portfolios to the dynamic portfolios.

		Static			***Long memory***			***Long memory vs. Static***	
Year	***μ (%)***	***σ (%)***	***SR***	***μ (%)***	***σ (%)***	***SR***	***p-value***	Δ_1	Δ_5
1994	5.874	2.054	0.921	9.968	5.108	1.173	0.680	398	353
1995	4.396	2.259	0.185	7.948	5.175	0.760	0.888	344	299
1996	5.958	2.100	0.944	9.170	5.205	0.997	0.593	310	263
1997	6.680	2.383	1.131	9.706	5.482	1.051	0.462	290	240
1998	6.012	2.717	0.743	9.121	5.772	0.896	0.644	298	244
1999	7.310	2.762	1.202	9.586	6.131	0.925	0.301	212	150
2000	7.062	2.873	1.068	9.741	6.369	0.913	0.409	251	184
2001	6.718	3.105	0.877	9.754	6.463	0.901	0.542	287	219
2002	5.718	3.256	0.529	9.983	6.761	0.886	0.715	408	334
2003	5.542	3.400	0.455	10.454	7.203	0.880	0.757	470	384
2004	5.492	3.351	0.447	10.507	7.566	0.843	0.767	477	378
2005	5.973	3.284	0.602	10.955	8.125	0.838	0.684	469	349
2006	6.060	3.220	0.641	13.439	9.301	0.973	0.698	697	530
2007	5.926	3.205	0.603	13.041	11.101	0.804	0.661	650	400
2008	5.321	3.389	0.391	14.957	15.036	0.716	0.714	846	369
2009	5.557	3.408	0.458	15.629	15.399	0.744	0.713	883	382

Table 6.18. Comparison of Rolling Window and Long Memory Volatility Timing

The table compares the out-of-sample performance of the rolling window and long memory volatility timing strategies. Expected returns are assumed constant and set to the unconditional mean of the estimation period. The rolling portfolio is constructed with rolling window unconditional covariance matrix estimator. The long memory volatility timing portfolio is constructed with the LM-EWMA model. The table reports the average annualised realised returns (μ), the annualised realised volatilities (σ), the Sharpe ratios (SR), the average annualised performance fees Δ_γ (in basis points) that an investor with a constant relative risk coefficient of γ is willing to pay to switch from the rolling portfolio to the long memory portfolio, and the corresponding breakeven transaction costs τ_γ (in basis points).

	Rolling			***Long memory***			***Long memory vs. Rolling***			
Rebalancing	***μ (%)***	***σ (%)***	***SR***	***μ (%)***	***σ (%)***	***SR***	Δ_1	Δ_5	τ_1	τ_5
Panel A. Stock-Bond Portfolio										
Daily	0.367	0.959	0.383	0.622	1.072	0.580	25	25	7	7
Weekly	0.364	0.890	0.409	0.579	1.016	0.570	21	21	13	13
Monthly	0.383	0.851	0.451	0.659	0.987	0.667	27	27	56	55
Panel B. S&P500-DJIA Portfolio										
Daily	0.133	0.438	0.303	0.135	0.466	0.291	0	0	0	0
Weekly	0.133	0.403	0.330	0.145	0.465	0.311	1	1	1	1
Monthly	0.137	0.404	0.339	0.165	0.465	0.355	3	3	9	9
Panel C. International Stock and Bond Portfolio										
Weekly	10.220	8.278	0.751	17.897	15.220	0.913	686	356	5	2
Monthly	9.870	7.031	0.835	15.938	11.537	1.035	565	390	7	5
Panel D. DJIA Portfolio										
Daily	4.195	2.297	0.085	4.981	5.807	0.169	64	7	2	0
Weekly	4.171	2.187	0.078	4.715	5.404	0.132	42	-7	2	–
Monthly	4.251	2.085	0.121	4.481	4.565	0.105	15	-19	3	–

Table 6.19. Portfolio Performance under the Long-Only Constraint

The table compares the out-of-sample performance of different optimal portfolios under the long-only constraint. Expected returns are assumed constant and set to the unconditional mean of the estimation period. For each strategy, the table reports the annualised average return μ (in percentage), the annualised volatility σ (in percentage), the Sharpe ratio (SR), the annualised performance fee Δ_1 (in basis points) that an investor with a constant relative risk coefficient of 1 is willing to pay to switch from the rolling portfolio to the long memory portfolio, and the corresponding breakeven transaction cost τ_1 (in basis points).

	Static			**GARCH-DCC vs. Static**					**LM-EWMA vs. Static**					**LM-EWMA-DCC vs. Static**				
	μ	σ	*SR*	μ	σ	*SR*	Δ_1	τ_1	μ	σ	*SR*	Δ_1	τ_1	μ	σ	*SR*	Δ_1	τ_1
Panel A. Stock-Bond Portfolio																		
Daily	0.27	0.87	0.31	0.36	0.92	0.39	9	2	0.59	1.06	0.56	32	9	0.51	0.96	0.53	24	5
Weekly	0.27	0.79	0.34	0.40	0.85	0.47	13	11	0.55	1.01	0.55	28	19	0.51	0.91	0.56	24	17
Monthly	0.27	0.74	0.37	0.45	0.80	0.56	17	40	0.64	0.99	0.65	36	75	0.58	0.86	0.67	31	71
Panel B. S&P500-DJIA Portfolio																		
Daily	0.08	0.52	0.15	0.10	0.43	0.24	2	3	0.14	0.44	0.30	6	5	0.14	0.44	0.31	6	6
Weekly	0.08	0.47	0.17	0.11	0.41	0.27	3	11	0.14	0.43	0.32	6	15	0.15	0.43	0.34	7	16
Monthly	0.08	0.46	0.17	0.12	0.42	0.29	4	30	0.14	0.43	0.34	7	45	0.15	0.42	0.35	7	47
Panel C. International Stock and Bond Portfolio																		
Weekly	4.09	2.39	0.04	4.27	2.04	0.13	19	20	4.53	2.11	0.25	45	35	4.48	1.83	0.26	41	72
Monthly	4.10	2.58	0.04	4.14	2.40	0.06	5	15	4.34	2.47	0.14	25	39	4.28	2.22	0.12	19	84
Panel D. DJIA Portfolio																		
Daily	4.20	1.72	0.12	4.23	2.14	0.11	1	0	4.01	2.44	0.01	-21	–	4.38	2.16	0.17	16	4
Weekly	4.21	1.59	0.13	4.35	2.13	0.16	13	9	4.09	2.36	0.04	-13	–	4.41	2.16	0.19	19	12
Monthly	4.21	1.54	0.14	4.42	1.93	0.22	20	46	4.31	2.03	0.15	10	15	4.55	1.97	0.28	33	75

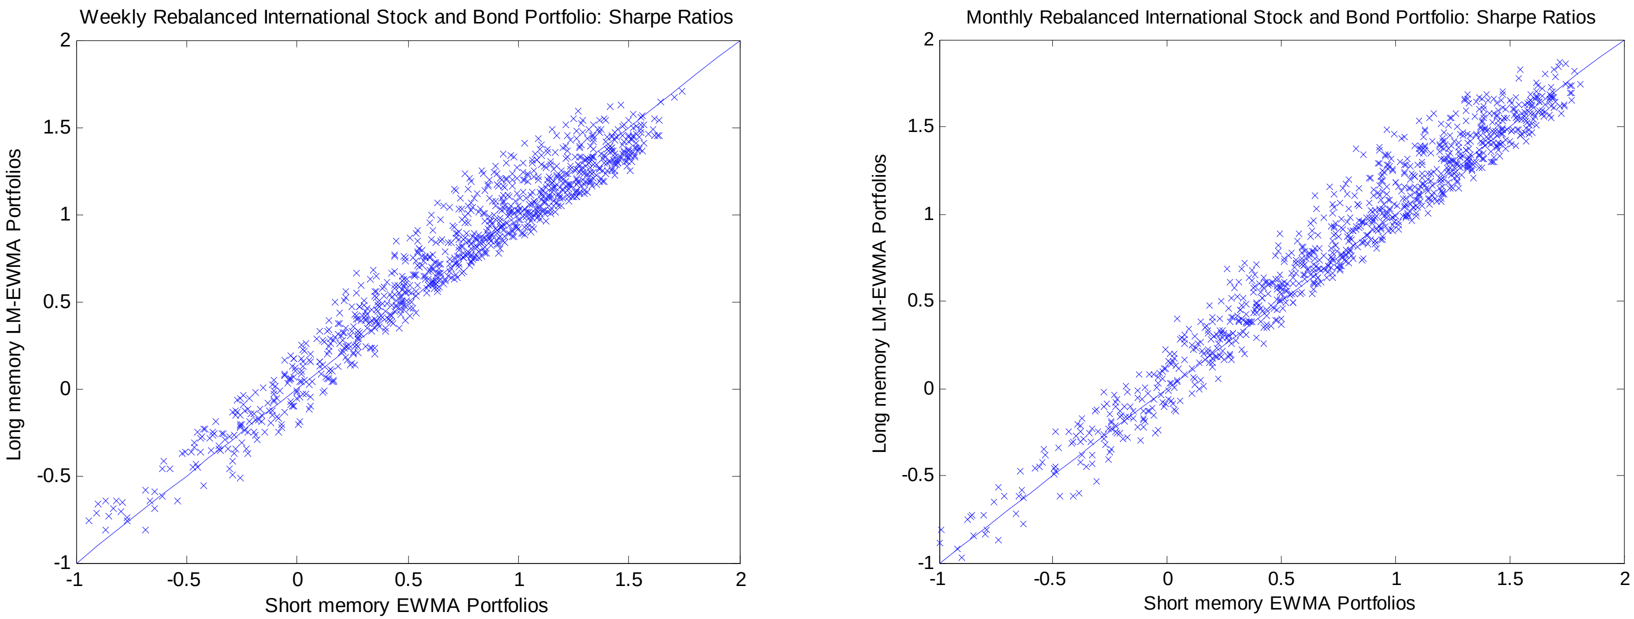

Figure 6.1. International Stock and Bond Portfolio: The Sharpe Ratios of the Short Memory and Long Memory Volatility Timing Strategies.

The figure plots the realised Sharpe ratios for 1,000 trials of the bootstrap experiment for the international stock and bond portfolio. Each dot represents a separate trial, plotting the realised Sharpe ratios for both the short memory and the long memory volatility timing portfolios.

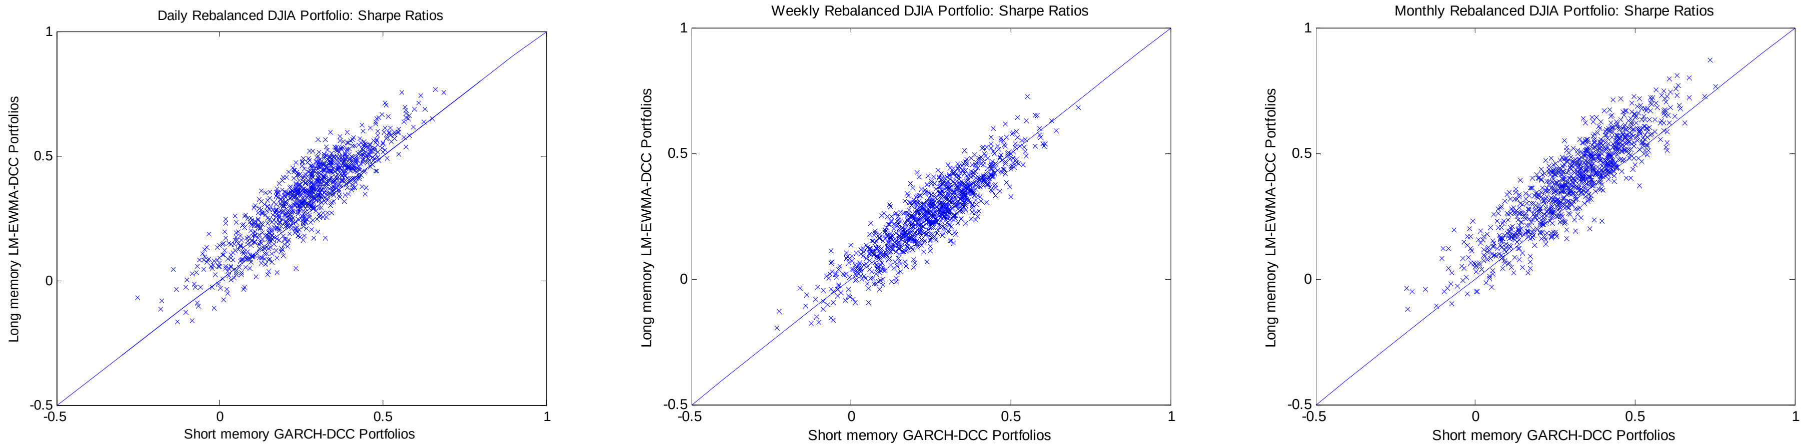

Figure 6.2. DJIA Portfolio: The Sharpe Ratios of the Short Memory and Long Memory Volatility Timing Strategies.

The figure plots the realised Sharpe ratios for 1,000 trials of the bootstrap experiment for the DJIA portfolio. Each dot represents a separate trial, plotting the realised Sharpe ratios for both the short memory and the long memory volatility timing portfolios.

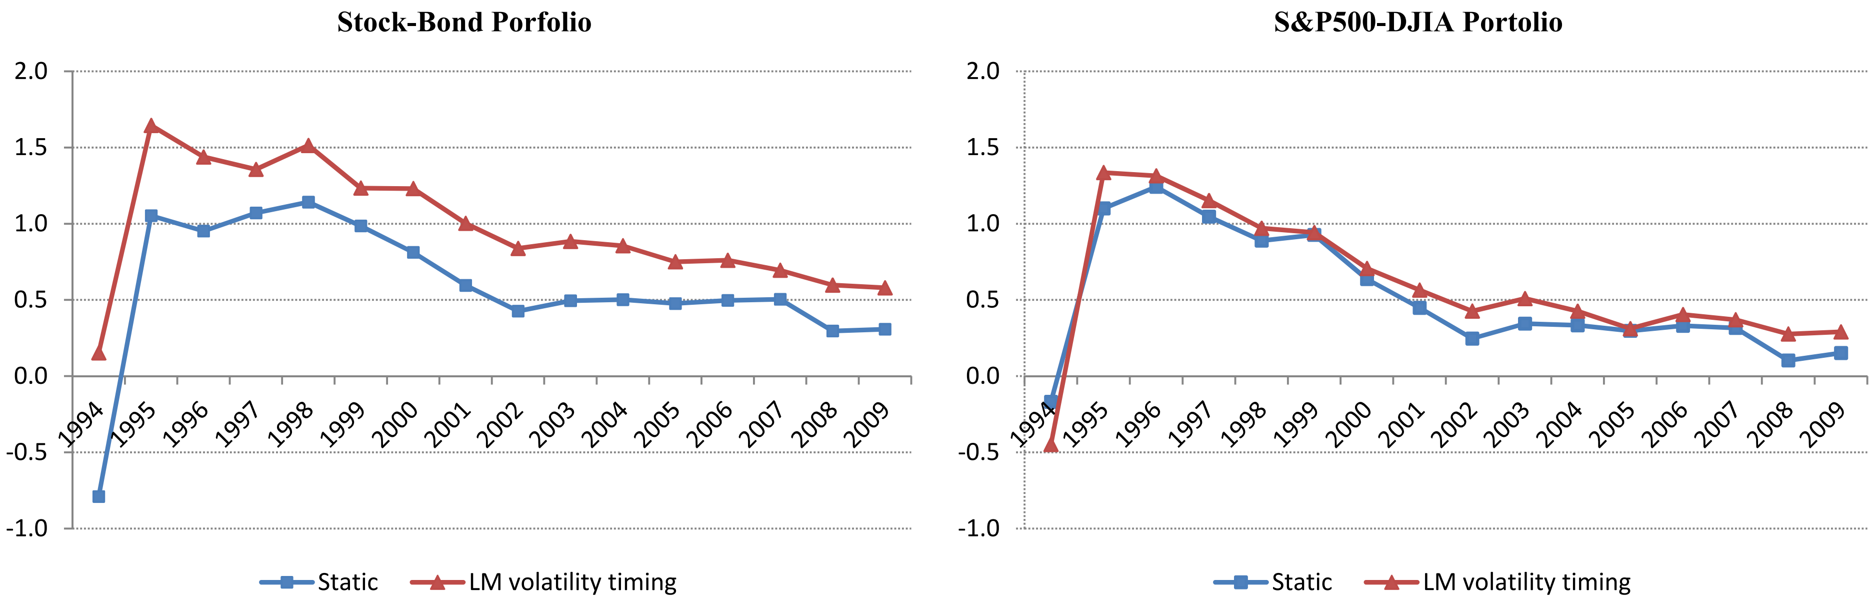

Figure 6.3. Year-on-Year Sharpe Ratios of the Static and Long Memory Volatility Timing Portfolios.

The figure plots the year-on-year performance of the static and long memory volatility timing portfolios, calculated over the period from the beginning of the testing period (January 2, 1994) to the end of the target year.

Chapter 7

Dynamic Factor Long Memory Conditional Volatility

Covariance matrix estimates are inevitably subject to estimation error. Estimation error may be excessive in the time-varying covariance matrix, not just because of the complexity in the estimation procedure of multivariate conditional volatility models, but also because of the shortened effective sample size due to exponential weighting of the volatility process. The high dimensionality and the use of the inverse covariance matrix typically encountered in asset allocation may further aggravate the estimation error problem. One popular approach to deal with estimation error is to impose a factor structure on the covariance matrix. Thanks to the ability of using some common factors to capture cross-sectional risk, factor models significantly reduce the number of parameters to be estimated, and hence reduce estimation error. Factor models also provide a better solution in the estimation of the inverse of the covariance matrix (see, for example, Fan et al., 2008 for a theoretical proof). The advantages of factor models relative to fully estimated covariance matrix estimators have been well documented in the literature and empirically confirmed in practice. Chan et al. (1999) study the performance of different fundamental factor models in a portfolio optimisation problem and show that factor models clearly improve forecasts of the covariance matrix. Similarly, Burmeister et al. (2003) consider a set of macroeconomic factors to construct superior portfolios. Commercially, MSCI BARRA has developed multifactor models covering the world's major equity markets. Recent studies incorporate the factor structure in the time-varying conditional volatility framework and suggest significant economic benefits. For example, Briner and Connor (2008) allow for the dynamic variations of returns' volatility and covariance by embedding an exponential weighting in the factor covariance matrix and prove that the conditional factor EWMA model outperforms the fully estimated EWMA model in terms of forecasts. Han (2006) develops a dynamic factor multivariate stochastic volatility model, which utilises unobserved factors to capture the dynamic behaviour of volatility (and also of returns). He shows that in the asset allocation framework, investors are better off in terms of utility when employing a dynamic factor model relative to using an unconditional covariance matrix estimator.

This chapter studies the benefits of imposing a factor structure in the long memory volatility dynamics to reduce estimation error in forecasts of the covariance matrix. In so doing, I first develop a dynamic factor long memory conditional volatility model that can be implemented in the context of the high dimensional covariance matrices typically encountered in risk management and asset allocation. Given the parsimony and outperformance of the long memory LM-EWMA and LM-EWMA-DCC models in the previous chapters, I employ the LM-EWMA model to capture the high persistence of financial asset volatility. The Orthogonal Factor Long Memory conditional volatility (OFLM) model is achieved by embedding the univariate long memory EWMA model of Zumbach (2006) into an orthogonal factor structure. I allow the new factor model to adopt richer specifications than normally assumed, in which both the factors and the idiosyncratic shocks are modelled with long memory behaviour in their volatilities. The OFLM model is a generalisation of the Factor Double ARCH model of Engle (2009). The factor-structured OFLM model is initially evaluated in terms of forecast performance against the fully estimated Long Memory Exponentially Weighted Moving Average model of Zumbach (2009b), using the procedure of Engle and Colacito (2006).[11] The performance of the OFLM model is also compared with that of a wide range of other multivariate conditional volatility models, both long memory and short memory, studied in the previous chapters. They include the three other long memory LM-EWMA-DCC, CGARCH-DCC and FIGARCH-DCC models and the two short memory Riskmetrics EWMA and GARCH-DCC models. The research then further evaluates the economic gains of employing the factor-structured long memory covariance matrix in the volatility timing framework of Fleming et al. (2001).[12] Portfolios constructed with the OFLM model are compared against those constructed with the other multivariate conditional volatility models, both short memory and long memory. All dynamic strategies are also evaluated against the static and the equally-weighted strategies. Other benchmarks include the traditional unconditional factor model, and the dynamic factor short memory EWMA and GARCH models.

As factor models are typically employed to improve estimates of the high dimensional covariance matrix, I employ the two multivariate portfolios for the empirical analysis, i.e., the international portfolio of 21 international stock indices and 13 international bond indices, and the US portfolio of the Dow Jones Industrial Index (DJIA) components. The analysis is conducted using data over the period 1 January 1988 to 31

[11] Details of the Engle and Colacito's (2006) framework have been given in Chapter 5.
[12] The volatility timing framework of Fleming et al. (2001) has been described in Chapter 6.

December 2009. Expected returns are assumed constant and investors periodically update their portfolios based on forecasts of the conditional covariance matrix. The results consistently show that the factor-structured OFLM model dominates the fully estimated LM-EWMA model at all forecast horizons, confirming the advantage of the factor structure to reduce estimation error. The OFLM model also generates impressive forecasts of the covariance matrix as compared to other multivariate conditional volatility benchmarks. Portfolios constructed with the dynamic factor long memory OFLM model also consistently dominate other static and dynamic portfolios. The factor structure also significantly reduces transaction costs, making the dynamic strategies more feasible in practice. The results apply to the two datasets, and are robust to estimation error in expected returns, the choice of risk aversion coefficient and the choice of estimation window. The dynamic factor long memory volatility models are also found to generally outperform the unconditional factor and the dynamic factor short memory volatility models.

The remainder of this chapter is organised as follows. Section 7.1 describes the Orthogonal Factor Long Memory conditional volatility model. Section 7.2 discusses the data and estimates the number of common factors. Section 7.3 evaluates the forecast performance of the new model, while Section 7.4 studies its economic usefulness in the volatility timing framework. The conclusion is given in Section 7.5.

7.1 The Orthogonal Factor Long Memory Conditional Volatility Model

Consider an n-dimensional vector of asset returns $\mathbf{r}_t = \left(r_{1t}, r_{2t}, \ldots, r_{nt}\right)'$. In the factor framework, asset returns are decomposed linearly into two parts, i.e., the part of returns that is correlated to a set of risk factors and the part of remaining asset-specific returns:

$$\mathbf{r}_t = \boldsymbol{\alpha} + \mathbf{B}\mathbf{f}_t + \boldsymbol{\varepsilon}_t \tag{7.1}$$

where $\mathbf{f}_t$ is a vector of q common risk factors $\mathbf{f}_t = \left(f_{1t}, f_{2t}, \ldots, f_{qt}\right)'$ $\left(q < n\right)$, $\mathbf{B}$ is an $n \times q$ matrix of factor loadings, and $\boldsymbol{\varepsilon}_t$ is a vector of asset-idiosyncratic returns. The vector of coefficients $\boldsymbol{\alpha}$ is set so that $E\left(\boldsymbol{\varepsilon}_t\right) = 0$. The idiosyncratic shocks are assumed to be uncorrelated with the factors and uncorrelated with each other. The conditional covariance matrix $\mathbf{H}_t$ can thus be represented as:

$$\mathbf{H}_t = \mathbf{B}\mathbf{\Omega}_t\mathbf{B}' + \mathbf{H}_{\varepsilon,t} \tag{7.2}$$

where $\mathbf{\Omega}_t$ is the covariance matrix of the common factors $\mathbf{f}_t$, and $\mathbf{H}_{\varepsilon,t}$ is the covariance matrix of the residuals $\boldsymbol{\varepsilon}_t$. The Orthogonal Factor Long Memory conditional volatility model assumes that the common factors are latent and orthogonal. The factor covariance matrix is then a diagonal matrix with the variance of $f_{i,t}$ on the i^{th} diagonal $\mathbf{\Omega}_t = diag\left\{\sigma^2_{f_{i,t}}\right\}$. Since the residuals are uncorrelated, $\mathbf{H}_{\varepsilon,t}$ is also a diagonal matrix $\mathbf{H}_{\varepsilon,t} = diag\left\{\sigma^2_{\varepsilon_{j,t}}\right\}$. For simplicity, I assume that the number of factors and the factor loadings are constant. Principal Components Analysis is applied to estimate the factor loadings **B** and to obtain the factors $\mathbf{f}_t$ and the residuals $\boldsymbol{\varepsilon}_t$.

In the OFLM model, both of the factors and the idiosyncratic shocks are time varying and exhibit long memory behaviour in volatility. The long memory volatilities of the factors and residuals are modelled using the univariate long memory EWMA model of Zumbach (2006). Remember that in the LM-EWMA model, long memory conditional volatility is defined as the weighted average of *K* standard (short memory) EWMA processes over increasing time horizons.

$$\sigma^2_{t+1} = \sum_{k=1}^{K} w_k \sigma^2_{k,t}, \tag{7.3}$$

where

$$\sigma^2_{k,t} = \mu_k \sigma^2_{k,t-1} + \left(1-\mu_k\right) r_t^2. \tag{7.4}$$

The EWMA process in (7.4) is characterised by its decay factor $\mu_k = \exp\left(-1/\tau_k\right)$ with the geometric characteristic time $\tau_k = \tau_1\left(\rho\right)^{k-1}$. Zumbach (2006) sets ρ to the value of $\sqrt{2}$. The long memory of the volatility process is embedded in the weights w_k, which are assumed to decay logarithmically $w_k = \frac{1}{C}\left(1-\frac{\ln\tau_k}{\ln\tau_0}\right)$, with the normalization factor *C* such that $\sum_k w_k = 1$. The long memory volatility can also be expressed in the form of the weighted average of past squared returns:

$$\sigma^2_{t+1} = \sum_{i=0}^{\infty} \lambda(i) r^2_{t-i} \tag{7.5}$$

with the logarithmically decaying weights $\lambda(i) = \sum_k w_k \left(1-\mu_k\right)\mu_k^i$. Applying the LM-EWMA model to the volatilities of the factors and of the residuals, I obtain:

$$\sigma^2_{f,t+1} = \sum_{i=0}^{\infty} \lambda(i) f^2_{t-i} , \tag{7.6}$$

and $$\sigma^2_{\varepsilon,t+1} = \sum_{i=0}^{\infty} \lambda(i) \varepsilon^2_{t-i} , \tag{7.7}$$

with $\lambda(i)$ defined above. Following Zumbach (2006), I set the optimal time parameter values at τ_0 = 1560 days = 6 years, τ_1 = 4 days and τ_K = 512 days, which is equivalent to $K = 15$.

The covariance matrices of the factors $\mathbf{\Omega}_t = diag\left\{\sigma^2_{f_{i,t}}\right\}$ and of the residuals $\mathbf{H}_{\varepsilon,t} = diag\left\{\sigma^2_{\varepsilon_{j,t}}\right\}$ are easily computed. Note that while the unconditional factors and residuals are uncorrelated, it does not necessarily imply that the conditional factors and residuals are uncorrelated. However, I assume no autocorrelation among the conditional factors and the conditional residuals so as to keep the parsimonious advantage of the model. $\mathbf{\Omega}_t$ and $\mathbf{H}_{\varepsilon,t}$ are then combined to estimate the conditional covariance matrix $\mathbf{H}_t$. Under the assumption of serially uncorrelated factors and residuals, the forecast of the conditional covariance matrix over h steps is given by

$$\mathbf{H}_{t+1:t+h} = \mathbf{B}\mathbf{\Omega}_{t+1:t+h}\mathbf{B}' + \mathbf{H}_{\varepsilon,t+1:t+h} . \tag{7.8}$$

Since the long memory volatility is the sum of different EWMA processes over increasing time horizons, the volatility forecasts are straightforward to obtain using a recursive procedure (see Chapter 5 for more details).

7.2 Data Analysis

To evaluate the economic benefits of incorporating a factor structure in the long memory volatility framework to estimate the high dimensional covariance matrix, I employ the two high dimensional portfolios, i.e., the international stock and bond portfolio and the DJIA portfolio. The descriptive summary of the two datasets has been given in Chapter 4. Again, the whole sample is divided into an estimation period and a forecast period. As previously, for the international stock and bond portfolio, the estimation period is from 1 Jan 1988 to 31 Dec 1993 (312 weekly observations) and the forecast period from 1 Jan 1994 to 31 Dec 2009 (835 observations). The estimation period of the DJIA portfolio ranges from 1 Mar 1990 to 29 Feb 1996 (1518 daily observations) and the forecast period from 1 Mar 1996 to 31 Dec 2009 (3483 observations). Expected returns are assumed constant and so the investor actively rebalances his dynamic portfolios periodically, based on changes in forecasts of the conditional covariance matrix. The estimation period is used to initiate the estimation of the conditional volatility models to generate out-of-sample forecasts of the covariance matrix. The forecasts are then used to compute optimal portfolio weights. Realised portfolio returns at the next step are calculated. Then the estimation window is rolled forward one step, models re-estimated, forecasts made, portfolios rebalanced and realised portfolio returns calculated, and so on until the end of the sample is reached. I estimate the conditional covariance matrix, using all seven multivariate volatility models for the DJIA portfolio. They include the factor long memory OFLM model, the four long memory LM-EWMA, LM-EWMA-DCC, CGARCH(1,1)-DCC and FIGARCH(1,*d*,1)-DCC models, and the two short memory Riskmetrics EWMA and GARCH(1,1)-DCC models. The FIGARCH-DCC model is, nevertheless, excluded in the international stock and bond portfolio owing to the prohibitively short estimation window.

I assume the number of factors constant and determine it by applying the test of Alessi et al. (2007). This is the generalisation of the information criterion (IC) of Bai and Ng (2002) to choose the number of factors in dynamic factor models. Bai and Ng's information criterion aims at minimizing the variances explained by the idiosyncratic components while penalizing the criterion to avoid over-parameterisation. Alessi et al. (2007) modify the penalty function in Bai and Ng's criterion so that the criterion is evaluated against a whole family of penalty functions rather than only one specific function as in Bai and Ng. I apply the Alessi et al. test to the estimation period. Figure 7.1 plots their modified Bai and Ng information criterion.[13] Visually, the number of factors corresponds to the second stable region, i.e., the plateau of the solid line associated with the second flat zero-level dashed line. The criterion identifies four common factors for the international stock and bond portfolio and three factors for the DJIA portfolio. However, to evaluate the sensitivity of the choice of the number of factors, I employ two, three, and four common factors in the empirical study.

[13] Figures are drawn employing the Matlab code provided by Barigozzi on his website at http://www.barigozzi.eu/research.html#codes.

7.3 Forecast Performance Evaluation

The dynamic factor long memory conditional volatility model is evaluated in terms of forecast performance against the other multivariate conditional volatility models employed in the previous chapters, using the asset allocation framework of Engle and Colacito (2006). These are the four long memory LM-EWMA, LM-EWMA-DCC, CGARCH(1,1)-DCC and FIGARCH(1,*d*,1)-DCC models and the two short memory Riskmetrics EWMA and GARCH(1,1)-DCC models. Remember that in the mean-variance optimisation framework, Engle and Colacito suggest the investor choose among competing covariance matrix forecasts based on the volatility of resulting portfolios. The best covariance matrix estimator produces the optimal portfolio with the lowest volatility, irrespective of both the expected returns and the target return. As with the previous experiments (see Section 5.3.2), I study the performance of the covariance matrix forecasts in two restricted cases, i.e., the global minimum variance portfolios where all expected returns are assumed to be equal, and the hedging portfolios where one asset is hedged against all other assets in the portfolio. In the hedging portfolios, expected returns are selected such that one entry is equal to one and all others are set to zero. The target excess return of the optimal portfolio is 1.[14]

Table 7.1 reports the out-of-sample volatilities of the global minimum variance strategy across different investment horizons. The 'Const' portfolio is the fixed weight portfolio constructed with the unconditional covariance matrix of the estimation period. The 'OFLMα' portfolio is constructed with the Orthogonal Factor Long Memory volatility model of α factors. Note that the FIGARCH-DCC model is excluded in the international stock and bond portfolio since its estimation requires a prohibitively high upper lag cut-off. Following standard practice in the literature, I use a truncation lag for the FIGARCH model of 1000. The lowest volatility is, again, normalised to 100. Consistent with the previous findings, the dynamic portfolios systematically dominate the constant portfolios in terms of low volatility at all rebalancing frequencies. Among the dynamic portfolios, the long memory volatility portfolios, especially those constructed with the two LM-EWMA models, consistently outperform the short memory volatility portfolios. The FIGARCH-DCC model, however, yields quite dismal results, probably owing to its complexity in estimation. Imposing a factor structure in the long memory covariance matrix brings further benefits. The factor-structured OFLM model generally

[14] Note, again, that the choice of the target return is immaterial in the sense that it does not affect the relative volatilities of portfolios.

gives superior forecasting results. In particular, the OFLM4 model produces the lowest portfolio volatility in three out of five experiments, and is among the top three models in the remaining two experiments. The results also suggest the importance of choosing the appropriate number of factors. Too few factors may not fully explain the comovement among assets, leading to the poor forecast performance of the OFLM2 model. Note also that though the information criterion identifies three common factors for the DJIA portfolio, the OFLM3 model underperforms the OFLM4 model, suggesting that the number of factors may change in the forecasting period.

The hedging portfolios bring mixed results. As the forecast performance of other multivariate conditional volatility models have been evaluated against each other in Chapter 5, in this section, I focus on the performance of the OFLM model. Table 7.2 reports results of the Diebold-Mariano joint tests for the equality of different models with all hedging vectors of expected returns. Each cell in the table corresponds to the test of the hypothesis that the OFLMα model in the row and the benchmark model in the column are equal in terms of volatility forecasting against the alternative that the OFLMα model is better or worse than the benchmark. A positive sign of the *t*-statistics means that the OFLM model is better than the benchmark and vice-versa. It is clearly demonstrated that the factor-structured OFLM model dominates the fully estimated long memory LM-EWMA and the short memory EWMA models in both versions of the Diebold-Mariano test. However, the OFLM model fails to dominate the DCC models in other experiments. The LM-EWMA-DCC model consistently generates better forecasts of the covariance matrix than the OFLM model, especially in the standard Diebold-Mariano tests. Though the equality hypothesis in terms of forecast performance cannot be rejected between the LM-EWMA-DCC and the OFLM models in the improved Diebold-Mariano tests, the sign of the *t*-statistics generally favours the LM-EWMA-DCC model. The comparison results between the OFLM model and the other two DCC models, the short memory GARCH-DCC and the long memory CGARCH-DCC models, are mixed. For example, while the improved Diebold-Mariano tests generally yield favourable results for the OFLM model, their standard version supports the GARCH-DCC model. Similarly, the OFLM model performs better at short forecast horizons, while the CGARCH-DCC model dominates at longer horizons.

7.4 The Economic Value of the Dynamic Factor Long Memory Volatility Timing Strategy

The economic value of the dynamic factor long memory volatility OFLM model is now analysed in greater depth, using the volatility timing framework of Fleming et al. (2001). Expected returns are, again, assumed constant and the investor periodically updates his portfolios based on forecasts of the conditional covariance matrix. Portfolio performance is evaluated using the out-of-sample Sharpe ratio, the abnormal return and the performance fee that the investor is willing to pay to switch from the static to the dynamic strategies. Transaction costs are also taken into consideration. To specifically evaluate the economic gains of imposing a factor structure in the long memory volatility framework, I initially compare the factor-structured OFLM model with the fully estimated long memory covariance matrix estimator, the LM-EWMA model. In Section 7.4.5, I extend the analysis to other benchmarks, considering all the other multivariate conditional volatility models studied so far. I also compare the factor long memory OFLM model with the traditional unconditional factor and the dynamic factor short memory conditional volatility models in Section 7.4.6.

7.4.1 Performance Analysis of the Dynamic Factor Long Memory Volatility Timing Strategy

Table 7.3 reports the out of sample performance of the international stock and bond portfolio with weekly and monthly rebalancing frequencies. The performance of the dynamic portfolios is compared with that of the ex ante static portfolio, constructed based on the sample mean and covariance matrix of the estimation period. Another benchmark is the equally weighted portfolio. It is, again, obvious that all dynamic portfolios outperform the static and the equally weighted portfolios. The long memory volatility models consistently produce portfolios with higher Sharpe ratios and positive abnormal returns. The passive investor is also willing to pay annualised performance fees of 338 up to 1242 *bps* to switch from the static to the dynamic long memory volatility timing strategies. Imposing the factor structure in the long memory covariance matrix brings further gains. The OFLM model with the recommended number of factors (four factors) generally produces portfolios with higher Sharpe ratios than those produced by the LM-EWMA model. Note that incorporating too few factors (two factors) may not explain enough movement in the covariance matrix, leading to poor results. Due to lower realised portfolio returns, the OFLM portfolio nevertheless generates lower performance fees than the LM-EWMA portfolio. With the relative risk aversion $\gamma = 1$, the performance fee of the OFLM4 portfolio is just around 420 *bps* compared to that of around 1100 *bps* of the LM-EWMA portfolio. However, the OFLM model yields much higher breakeven transaction costs, which make it much easier to be implemented in practice. For example, a week trader with $\gamma = 1$ is only better off with the LM-EWMA portfolio if his realised transaction cost is lower than 7 *bps*, compared to that of 34 *bps* if he employs the corresponding OFLM4 portfolio. As expected, less frequent trading yields higher transaction costs. The breakeven transaction costs for a month trader are much higher than those for a week trader, making dynamic trading more feasible. The breakeven transaction costs τ_γ, by construction, can be interpreted as the performance fees after taking transactions into account, i.e., the transaction-adjusted performance fees. Note that unlike the unadjusted performance fees, which are the absolute fees and reported in annualised basis points, the transaction-adjusted performance fees are the relative percentage of the value traded and reported in basis points at the rebalancing frequency. In this sense, when trading costs are allowed for, the investor is willing to pay more to switch from the static portfolios to the OFLM portfolio than to the LM-EWMA portfolio.

Similar results are identified with the DJIA portfolio (Table 7.4). The dynamic OFLM strategy, especially the recommended one with three factors, consistently dominates other strategies, including the static, equally weighted and dynamic LM-EWMA portfolios in all performance measures and rebalancing frequencies. For example, the investor who rebalances his portfolio monthly is willing to pay performance fees from 55 *bps* to 75 *bps*, or from 35 *bps* to 50 *bps* of the value traded when adjusted for trading, to switch from the static strategy to the dynamic OFLM strategy. The outperformance of the OFLM model relative to the LM-EWMA model is more clearly marked with this dataset. The Sharpe ratios of the OFLM portfolios with three and four factors are generally twice as much as those of the LM-EWMA portfolio. The OFLM portfolios also dominate in terms of abnormal returns, performance fees and breakeven transaction costs. Especially, the low breakeven transaction costs of the LM-EWMA portfolio make them undesirable for traders, who are, on the contrary, willing to adopt the OFLM portfolios. Again, it is of extreme importance to include the appropriate number of factors in estimating the covariance matrix (see the poor performance of the OFLM2 model). The results also suggest that incorporating not enough factors is more detrimental to portfolio performance than incorporating more factors.

7.4.2 Estimation Error in Expected Returns

Again, I follow Fleming et al.'s (2001) recommendation to consider a range of expected returns generated via a block bootstrap procedure. An artificial sample of 4000 observations is created by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. I then estimate the unconditional mean and covariance matrix of these artificial return series. Dynamic portfolios are constructed from the bootstrap unconditional expected returns and forecasts of the conditional covariance matrix, while static portfolios employ both bootstrap mean and covariance matrix estimates. I repeat this procedure 1000 times, studying the economic gains of volatility timing across a wide range of plausible vectors of expected returns.

Table 7.5 reports the average results across the 1000 bootstrap vectors of expected returns for the international stock and bond portfolio. Again, it is clear that the investor is better off switching from the static to the dynamic, and from the fully estimated LM-EWMA to the factor-structured OFLM volatility timing strategies. The OFLM model generates portfolios with positive abnormal returns and performance fees, and higher Sharpe ratios than those produced by the static unconditional covariance matrix estimator in all trials. The OFLM portfolio also dominates the LM-EWMA portfolio in terms of Sharpe ratios, abnormal returns and breakeven transaction costs (or trading-adjusted performance fees), though the LM-EWMA portfolio outperforms in terms of unadjusted performance fees due to their higher realised portfolio returns. Among the factor models, the OFLM4 model performs best. For example, the OFLM4 model produces weekly rebalanced portfolios with the average Sharpe ratio of 1.537, as compared to that of just 0.744 generated by the LM-EWMA model. Figure 7.2 illustrates the outperformance of the OFLM4 relative to the LM-EWMA portfolios in terms of Sharpe ratios and adjusted-performance fees (for $\gamma = 1$). Similar results are reported for the DJIA portfolio, where the OFLM model generates the best performing portfolios in terms of all performance measures and across all investment horizons (Table 7.6).

7.4.3 Sensitivity to Risk Aversion Coefficient

Different risk aversion coefficients λ are now employed to check the robustness of the dynamic portfolio performance. So far all reported results are based on $\lambda = 1$. For each value of risk aversion coefficients λ, I repeat the experiment with 1,000 bootstrap vectors of expected returns. Table 7.7 compares the static and dynamic long memory strategies across different risk aversion coefficients for the international stock and bond portfolio. To save space, I only report the results of the factor-structured OFLM4 model and the unstructured LM-EWMA model. Expectedly, when the investor is more risk averse (higher values of λ), he will choose more conservative portfolios with lower risk and lower expected returns. The Sharpe ratios are approximately the same across all risk aversion levels, with the slight difference due to the bootstrap procedure. Consistent with the previous results, the dynamic portfolios, especially the factor portfolios, generate higher Sharpe ratios than the static portfolios in most bootstrap trials. The average Sharpe ratios of the factor portfolios are even three times as much as those of the static portfolios. Investors are also willing to pay to switch from the static portfolios to the dynamic portfolios. For example, the investor is willing to pay annualised performance fees of around 50 to 270 *bps* to switch from the static portfolios to the OFLM4 portfolios. Compared to the fully estimated LM-EWMA model, the factor OFLM model produces portfolios with higher Sharpe ratios, lower performance fees due to lower realised portfolio returns, but higher transaction-adjusted performance fees due to less trading. The OFLM model is hence more applicable and easier to be implemented in practice. Similar results apply to the DJIA portfolio where the OFLM portfolio dominates across all risk aversion coefficients and investment horizons (Table 7.8).

7.4.4 Sensitivity to Estimation Window

As the factor loadings may be sensitive to the estimation window, different estimation windows are used for another robustness check. Experiments are, again, conducted with the bootstrap vectors of expected returns. Figure 7.3 shows the average Sharpe ratios of the dynamic strategies for different estimation windows. The OFLM model generally dominates the LM-EWMA model for all different estimation windows (4, 6, 8 and 10 years of weekly and daily data). It is worth noticing that the Sharpe ratios of the factor model tend to decline with long windows (10 years of data), suggesting that estimating the factor loadings with too distant information may be inaccurate.

7.4.5 Performance Analysis with Other Multivariate Conditional Volatility Benchmarks

The OFLM model is now evaluated against other multivariate conditional volatility models, including the two short memory Riskmetrics EWMA and GARCH(1,1)-DCC models and the three other long memory LM-EWMA-DCC, FIGARCH(1,*d*,1)-DCC and CGARCH(1,1)-DCC models. Table 7.9 compares the out-of-sample performance of the OFLM strategy with that of other volatility timing strategies for the DJIA portfolio. Results of the international stock and bond portfolio are similar and are reported in Table 7.10. It is clearly shown that the portfolios constructed using the OFLM model generally outperform those constructed using other conditional volatility models with all performance measures and across all investment horizons. The OFLM model produces portfolios with higher Sharpe ratios than those produced by any other conditional volatility model in around 72% – 99% of all trials. Also, as the dynamic factor portfolios require less trading than other conditional volatility portfolios, the investor is always better off in terms of transaction costs when implementing the dynamic factor model. The breakeven transaction cost of Han (2006) is not applicable here as the dynamic factor portfolios typically have fewer transactions than the benchmarks. To get a sense of the amount of trading required to implement each strategy, I report the portfolio *turnover* instead. Turnover is defined as the average value traded for all stocks in each period, and equal to $\frac{1}{T}\sum\left|w_{t+1}-\frac{w_t\left(1+r_{p,t+1}\right)}{w_t\left(r_{p,t+1}-r_f\right)+r_f+1}\right|$. As expected, the dynamic factor model consistently has the lowest turnover. Note again that among the conditional volatility models, the long memory LM-EWMA-DCC and CGARCH-DCC models perform better than the two short memory EWMA and GARCH-DCC models, suggesting the benefits of allowing for long memory in volatility modelling. Though the FIGARCH model is also designed to capture long memory volatility, its high degree of parameterisation evidently hinders its performance, making it the worst model.

7.4.6 Performance Analysis with Other Factor Models

Other benchmarks are used to evaluate the performance of the dynamic factor long memory model. Previously, I examine the gains from embedding a factor structure in the long memory volatility framework. I now approach from the opposite direction, evaluating the benefits of allowing for long memory volatility dynamics in the factor structure. In so doing, I compare the performance of the long memory factor model with the traditional statistical factor model. The OFLMα portfolio is hence evaluated against the corresponding traditional α-factor portfolio, where α is the number of factors. Again, the bootstrap procedure is employed to account for estimation error in expected returns. Results are reported in Table 7.11. The long memory factor model generally dominates the traditional factor model. In particular, in the international stock and bond dataset, the OFLM portfolios consistently yield higher Sharpe ratios in most of the trials. The investor is hence willing to pay up to 40 *bps* to switch from the traditional α-factor strategy to the long memory α-factor strategy. The OFLM portfolios perform a bit worse with the DJIA dataset, however the OFLM3 and OFLM4 portfolios still generate higher Sharpe ratios, higher performance fees and higher breakeven transaction costs than the corresponding traditional factor portfolios. The performance of the OFLM model is also compared with that of the orthogonal factor short memory volatility models, in which the volatilities of the factors and the residuals follow the EWMA and GARCH processes. The results are similar. The short memory factor models are found to consistently outperform the unconditional factor models, but they are generally dominated by the long memory factor models (see Appendices 7.1 and 7.2).

7.5 Conclusion

The chapter develops a dynamic factor long memory conditional volatility (OFLM) model that combines the long memory behaviour of volatility with the factor structure. The new model can capture the highly persistent property of financial volatilities observed in practice, while reducing estimation error in modelling high dimensional covariance matrices. The factor-structured OFLM model generally produces forecasts of the covariance matrix that are more economically useful than those produced by other multivariate conditional volatility models, both short memory and long memory. In the volatility timing framework, portfolios constructed with the OFLM model also dominate the static and other dynamic volatility timing portfolios for all rebalancing frequencies of up to one month. Employing the factor structure significantly reduces transaction costs, making dynamic trading more feasible. The findings also suggest that combining long memory volatility dynamics and a factor structure yields better results than employing long memory volatility or a factor structure alone. The results apply to both high dimensional datasets and are robust to estimation error in expected returns, the choice of risk aversion coefficient, and the length of estimation window.

Factor loadings are assumed constant in the OFLM model. It would be of interest to relax this assumption, developing a dynamic factor model with conditional betas. Also, expected returns may be time-varying and another promising line of research is to construct dynamic autoregressive factor models that are able to estimate both expected returns and the covariance matrix for asset allocation.

Table 7.1. Comparison of Out-of-Sample Volatilities

The table reports the out-of-sample volatilities of the global minimum variance portfolios across different investment horizons. The 'OFLMα' portfolio is constructed with the Orthogonal Factor Long Memory volatility model with α factors. The FIGARCH-DCC model is excluded in the international stock and bond portfolio owing to short estimation window. Daily, weekly, monthly are the rebalancing frequencies. The lowest volatility in each column is normalised to 100.

	International Stock and Bond Portfolio		**DJIA Portfolio**		
	Weekly	*Monthly*	*Daily*	*Weekly*	*Monthly*
Const	109.661	107.906	126.391	114.250	123.192
EWMA	135.112	105.112	126.670	129.893	111.162
GARCH-DCC	121.861	118.310	103.958	105.993	108.153
LM-EWMA	109.333	100.000	104.921	107.466	104.375
LM-EWMA-DCC	113.021	103.162	100.000	103.825	103.336
CGARCH-DCC	112.924	111.517	101.066	102.086	104.115
FIGARCH-DCC	–	–	125.296	117.817	122.747
OFLM2	102.241	106.403	104.943	107.021	105.204
OFLM3	101.162	107.183	104.123	103.427	103.895
OFLM4	100.000	103.248	100.887	100.000	100.000

Table 7.2. Diebold–Mariano Tests of the Hedging Portfolios

The table reports the *t*-statistics of the Diebold–Mariano joint tests for the hedging multivariate portfolios, using the improved test of Engle and Colacito (2006). Panel A corresponds to the international stock and bond portfolio, while Panel B corresponds to the DJIA portfolio. The *t*-statistics for the standard test are reported in parentheses. The 'OFLMα' portfolio is constructed with the Orthogonal Factor Long Memory volatility model with α factors. The FIGARCH-DCC model is excluded in the international stock and bond portfolio owing to short estimation window. A positive number indicates that the model in the row is better than the model in the column, and vice-versa. *, ** and *** denote rejection of the equality hypothesis of the two models at 10%, 5% and 1% significance level.

Panel A. International Stock and Bond Portfolio								
	EWMA	**GARCH DCC**	**LM-EWMA**	**LM-EWMA DCC**	**CGARCH DCC**	**OFLM2**	**OFLM3**	**OFLM4**
A1. Weekly Rebalancing								
OFLM2	2.84*** (5.70***)	1.07 (-0.41)	2.06** (3.19***)	-0.57 (-4.28***)	0.98 (0.14)	–	(-3.17***) (-5.91***)	-2.08** (-6.51***)
OFLM3	3.20*** (6.98***)	1.57 (-0.94)	2.47** (4.50***)	0.04 (-2.80***)	1.55 (1.42)	(3.17***) (5.91***)	–	-0.59 (-3.81***)
OFLM4	3.56*** (7.60***)	1.75* (1.69*)	2.80*** 5.16***)	0.23 (-1.78*)	1.73* (2.12**)	2.08** (6.51***)	0.59 (3.81***)	–
A2. Monthly Rebalancing								
OFLM2	1.33 2.85***	-0.15 (-0.75)	0.71 (0.38)	-1.73* (-3.45***)	-1.04 (-2.13**)	–	-1.45 (-2.25**)	-1.90* (-3.42***)
OFLM3	1.62 (3.74***)	0.18 (-0.25)	0.95 (1.10)	-1.62 (-2.97***)	-0.85 (-1.61)	1.45 (2.25**)	–	-1.72* (-3.05***)
OFLM4	1.95* (5.11***)	0.66 (0.53)	1.25 (2.30**)	-1.39 (-2.32**)	-0.56 (-0.75)	1.90* (3.42***)	1.72* (3.05***)	–

Panel B. DJIA Portfolio									
	EWMA	**GARCH DCC**	**LM EWMA**	**LM-EWMA DCC**	**CGARCH DCC**	**FIGARCH DCC**	**OFLM2**	**OFLM3**	**OFLM4**
B1. Daily Rebalancing									
OFLM2	9.03*** (30.26***)	-1.4248 (-11.51***)	7.04*** (15.37***)	-2.75*** (-12.41***)	-1.83* (-11.22***)	3.13*** (5.98***)	–	-4.65*** (-9.88***)	-3.71*** (-9.68***)
OFLM3	9.22*** (30.74***)	-0.8475 (-9.18***)	7.95*** (17.39***)	-2.28** (-10.35***)	-1.24 (-8.89***)	3.11*** (6.56***)	4.65*** (9.88***)	–	-2.87*** (-6.45***)
OFLM4	9.06*** (30.55***)	0.8621 (-7.69***)	8.07*** (19.77***)	-1.0816 (-10.42***)	0.67 (-7.36***)	3.49*** (6.95***)	3.71*** (9.68***)	2.87*** (6.45***)	–
B2. Weekly Rebalancing									
OFLM2	5.42*** (11.35***)	-0.10 (-8.48***)	4.48*** (7.97***)	-1.61 (-9.51***)	-0.13 (-6.35***)	1.37 (2.19**)	–	-3.33*** (-6.22***)	-2.37** (-7.32***)
OFLM3	5.46*** (12.06***)	0.78 (-6.27***)	4.67*** (9.00***)	-0.93 (-7.77***)	0.43 (-4.20***)	1.38 (2.35**)	3.33*** (6.22***)	–	-1.36 (-4.95***)
OFLM4	6.09*** (12.95***)	1.22 (-3.63***)	5.07*** (10.02***)	-0.18 (-5.45***)	1.13 (-2.42**)	1.45 (2.47**)	2.37** (7.32***)	1.36 (4.95***)	–
B3. Monthly Rebalancing									
OFLM2	4.96*** (7.67***)	-0.21 (-5.36***)	2.48** (2.87***)	-2.43** (-8.36***)	-3.19*** (-6.62***)	1.42 (1.66*)	–	-2.36** (-3.19***)	-2.36** (-4.34***)
OFLM3	5.48*** (8.48***)	0.12 (-3.85***)	3.14*** (3.88***)	-1.45 (-7.08***)	-2.34** (-5.46***)	1.41 (1.80*)	2.36** (3.19***)	–	-1.42 (-3.02***)
OFLM4	(5.90***) (9.31***)	0.24 (-2.41**)	3.92*** (5.02***)	-0.79 (-4.90***)	-1.83* (-4.25***)	1.48 (1.94*)	2.36** (4.34***)	1.42 (3.02***)	–

Table 7.3. Portfolio Performance of the International Stock and Bond Portfolio

The table compares the out-of-sample performance of the optimal international stock and bond portfolio. Panels A and B report results for the weekly and monthly rebalanced portfolios, respectively. 1/N is the equally weighted portfolio. The static portfolio is constructed using the constant mean and covariance matrix of the estimation period. The 'OFLMα' portfolio is constructed with the Orthogonal Factor Long Memory volatility model with α factors. For each dynamic volatility timing strategy, the table reports the annualised average return (μ), the annualised volatility (σ), the Sharpe ratio (SR), the annualised abnormal return (M2) to the static portfolio, the annualised performance fee (in basis points) Δ_γ that an investor with a constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio, and the breakeven transaction cost τ_γ (in basis points at the rebalancing frequency) that he is better off with the dynamic strategy.

	μ (%)	*σ (%)*	***SR***	***M2 (%)***	Δ_1	Δ_5	τ_1	τ_5
Panel A. Weekly rebalancing								
1/N	3.911	12.996	-0.007					
Static	4.393	3.836	0.103					
Volatility timing strategies								
LM-EWMA	17.897	15.220	0.913	3.109	1242	805	7	4
OFLM2	7.853	4.217	0.914	3.112	344	338	33	31
OFLM3	7.932	4.088	0.962	3.296	353	349	30	30
OFLM4	8.622	4.410	1.048	3.628	421	411	34	32
Panel B. Monthly rebalancing								
1/N	4.071	14.064	0.005					
Static	4.354	3.874	0.092					
Volatility timing strategies								
LM-EWMA	15.938	11.537	1.035	3.655	1099	858	12	9
OFLM2	7.903	4.435	0.880	3.055	353	343	64	61
OFLM3	8.026	4.141	0.972	3.412	366	362	60	59
OFLM4	8.701	4.504	1.044	3.690	432	421	67	64

Table 7.4. Portfolio Performance of the DJIA Portfolio

The table compares the out-of-sample performance of the optimal DJIA portfolios. Panels A, B and C report results for the daily, weekly and monthly rebalanced portfolios, respectively. 1/N is the equally weighted portfolio. The static portfolio is constructed using the constant mean and covariance matrix of the estimation period. The 'OFLMα' portfolio is constructed with the Orthogonal Factor Long Memory volatility model with α factors. For each dynamic volatility timing strategy, the table reports the annualised average return (μ), the annualised volatility (σ), the Sharpe ratio (SR), the annualised abnormal return (M2) to the static portfolio, the annualised performance fee (in basis points) Δ_γ that an investor with a constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio, and the breakeven transaction cost τ_γ (in basis points at the rebalancing frequency) that he is better off with the dynamic strategy.

	μ (%)	σ (%)	SR	M2 (%)	Δ_1	Δ_5	τ_1	τ_5
Panel A. Daily rebalancing								
1/N	5.428	21.191	0.067					
Static	4.166	1.796	0.092					
Volatility timing strategies								
LM-EWMA	4.981	5.807	0.169	0.137	66	5	2	0
OFLM2	4.328	3.721	0.088	-0.008	11	-10	1	–
OFLM3	4.974	3.408	0.286	0.347	77	60	7	6
OFLM4	5.011	3.435	0.294	0.363	80	63	8	6
Panel B. Weekly rebalancing								
1/N	5.509	19.265	0.078					
Static	4.170	1.660	0.102					
Volatility timing strategies								
LM-EWMA	4.715	5.404	0.132	0.050	41	-12	2	–
OFLM2	4.112	3.823	0.029	-0.121	-12	-36	–	–
OFLM3	4.753	3.484	0.216	0.189	54	35	11	7
OFLM4	4.625	3.371	0.185	0.138	41	24	8	5
Panel C. Monthly rebalancing								
1/N	5.427	19.835	0.072					
Static	4.182	1.649	0.111					
Volatility timing strategies								
LM-EWMA	4.481	4.565	0.105	-0.009	21	-16	4	–
OFLM2	4.439	3.909	0.112	0.003	19	-6	12	–
OFLM3	5.000	3.431	0.291	0.298	77	59	49	38
OFLM4	4.998	3.616	0.276	0.273	76	55	48	35

Table 7.5. Average Portfolio Performance of the International Stock and Bond Portfolio with Bootstrap Experiments

The table compares the average out-of-sample performance of the optimal international stock and bond portfolio across a wide range of expected returns. A bootstrap procedure is applied to control for estimation error in expected returns. I generate an artificial sample of 4000 observations by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. The procedure is repeated with 1,000 trials. Panels A and B report results for the weekly and monthly rebalanced portfolios, respectively. The static portfolios are constructed using the bootstrap unconditional means and covariance matrices. The 'OFLMα' portfolio is constructed with the Orthogonal Factor Long Memory volatility model with α factors. For each dynamic strategy, the table reports the annualised average return (μ), the annualised average volatility (σ), the average Sharpe ratio (SR), the *p*-value (proportion) that the dynamic strategy outperforms the static alternative in terms of the Sharpe ratio, the average abnormal return to the static portfolio (M2), the average annualised performance fee (in basis points) Δ_γ that an investor with a constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio, and the average breakeven transaction cost τ_γ (in basis points at the rebalancing frequency) that he is better off with the dynamic strategy.

	μ (%)	*σ (%)*	*SR*	*p-value*	*M2 (%)*	Δ_1	Δ_5	τ_1	τ_5
Panel A. Weekly rebalancing									
Static	5.557	3.408	0.458						
Volatility timing strategies									
LM-EWMA	15.629	15.399	0.744	0.713	0.988	883	382	4.7	2.3
OFLM2	7.247	2.522	1.291	1.000	2.855	172	182	17.9	19.0
OFLM3	7.486	2.611	1.338	1.000	3.015	195	205	18.4	19.4
OFLM4	8.096	2.670	1.537	1.000	3.697	256	265	22.9	23.7
Panel B. Monthly rebalancing									
Static	5.507	3.557	0.426						
Volatility timing strategies									
LM-EWMA	14.606	12.713	0.832	0.740	1.463	827	486	8.4	5.2
OFLM2	7.214	2.515	1.285	1.000	3.069	174	187	36.6	39.4
OFLM3	7.478	2.542	1.374	1.000	3.388	200	213	37.7	40.1
OFLM4	8.069	2.577	1.584	1.000	4.137	259	272	46.7	49.0

Table 7.6. Average Portfolio Performance of the DJIA Portfolio with Bootstrap Experiments

The table compares the average out-of-sample performance of the optimal DJIA portfolio across a wide range of expected returns. A bootstrap procedure is applied to control for estimation error in expected returns. I generate an artificial sample of 4000 observations by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. The procedure is repeated with 1,000 trials. Panels A, B and C report results for the daily, weekly and monthly rebalanced portfolios, respectively. The static portfolios are constructed using the bootstrap unconditional means and covariance matrices. The 'OFLMα' portfolio is constructed with the Orthogonal Factor Long Memory volatility model with α factors. For each dynamic strategy, the table reports the annualised average return (μ), the annualised average volatility (σ), the average Sharpe ratio (SR), the *p*-value (proportion) that the dynamic strategy outperforms the static alternative in terms of Sharpe ratios, the average abnormal return to the static portfolio (M2), the average annualised performance fees (in basis points) Δ_γ that an investor with a constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio, and the average breakeven transaction costs τ_γ (in basis points at the rebalancing frequency) that he is better off with the dynamic strategy.

	μ (%)	*σ (%)*	*SR*	*p-value*	*M2 (%)*	Δ_1	Δ_5	τ_1	τ_5
Panel A. Daily rebalancing									
Static	4.620	3.197	0.201						
Volatility timing strategies									
LM-EWMA	5.489	5.080	0.290	0.726	0.307	79	48	2.1	1.3
OFLM2	4.864	2.774	0.307	0.788	0.356	26	32	2.6	3.3
OFLM3	5.193	2.840	0.414	0.914	0.697	59	64	5.6	6.2
OFLM4	5.318	2.984	0.434	0.930	0.762	71	74	6.3	6.6
Panel B. Weekly rebalancing									
Static	4.634	2.856	0.230						
Volatility timing strategies									
LM-EWMA	5.270	4.739	0.265	0.592	0.118	56	28	3.4	1.7
OFLM2	4.787	2.786	0.280	0.647	0.158	16	17	3.3	3.7
OFLM3	5.132	2.822	0.396	0.873	0.488	50	51	10.4	10.7
OFLM4	5.155	2.926	0.389	0.856	0.465	52	52	10.1	10.1
Panel C. Monthly rebalancing									
Static	4.629	2.815	0.234						
Volatility timing strategies									
LM-EWMA	5.140	3.772	0.302	0.664	0.219	48	35	9.6	7.2
OFLM2	4.870	2.617	0.333	0.758	0.303	25	27	16.6	19.1
OFLM3	5.171	2.648	0.439	0.904	0.599	55	57	35.4	37.5
OFLM4	5.251	2.799	0.442	0.906	0.606	62	63	37.7	38.7

Table 7.7. Comparison of the Volatility Timing and Static Strategies Using Different Risk Aversion Coefficients: International Stock and Bond Portfolio

The table compares the average out-of-sample performance of the static and dynamic strategies for the international stock and bond portfolio using different risk aversion coefficients λ. A bootstrap procedure is applied to control for estimation error in expected returns. I generate an artificial sample of 4000 observations by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. The procedure is repeated with 1,000 trials. Panels A and B report results for the weekly and monthly rebalanced portfolios, respectively. The static portfolios are constructed using the bootstrap unconditional means and covariance matrices. The table reports the annualised average returns μ (in percentage), the annualised volatilities σ (in percentage), the Sharpe ratios (SR), the p-values (proportions) that the dynamic strategies outperform the static alternatives in terms of Sharpe ratios, the average annualised performance fees Δ_γ (in basis points) that an investor with a constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolios, and the average breakeven transaction costs τ_γ (in basis points at the rebalancing frequency) that he is better off with the dynamic strategies.

	Static			***Long memory***			***Long memory vs. Static***					***OFLM4***			***OFLM4 vs. Static***				
λ	μ	σ	***SR***	μ	σ	***SR***	***p-value***	Δ_1	Δ_5	τ_1	τ_5	μ	σ	***SR***	***p-value***	Δ_1	Δ_5	τ_1	τ_5
Panel A. Weekly rebalancing																			
1	5.56	3.41	0.46	15.63	15.40	0.74	0.71	883	382	4.7	2.3	8.10	2.67	1.54	1.00	256	265	22.9	23.7
2	4.78	1.71	0.46	10.03	7.71	0.77	0.73	493	367	5.4	4.1	6.07	1.34	1.55	1.00	129	132	23.1	23.5
3	4.52	1.14	0.46	8.00	5.16	0.76	0.73	334	277	5.6	4.7	5.37	0.89	1.55	1.00	85	86	22.8	23.1
4	4.39	0.85	0.46	6.91	3.89	0.74	0.73	244	212	5.5	4.8	5.03	0.67	1.54	1.00	64	65	22.7	22.9
5	4.31	0.68	0.46	6.41	3.11	0.77	0.74	204	184	5.8	5.2	4.82	0.53	1.55	1.00	51	52	22.9	23.1
Panel B. Monthly rebalancing																			
1	5.51	3.56	0.43	16.99	12.71	0.83	0.74	827	486	8.4	5.2	8.07	2.58	1.58	1.00	259	272	46.7	49.0
2	4.76	1.79	0.43	10.76	6.32	0.86	0.76	454	370	9.4	7.7	6.06	1.29	1.61	1.00	131	134	47.0	48.2
3	4.51	1.19	0.43	8.46	4.22	0.86	0.76	305	268	9.5	8.4	5.36	0.86	1.60	1.00	86	88	46.6	47.3
4	4.38	0.89	0.43	7.27	3.21	0.83	0.75	222	200	9.3	8.5	5.02	0.64	1.59	1.00	65	65	46.3	46.9
5	4.30	0.71	0.43	6.69	2.57	0.86	0.76	186	172	9.9	9.2	4.82	0.51	1.60	1.00	52	52	46.7	47.2

Table 7.8. Comparison of the Volatility Timing and Static Strategies Using Different Risk Aversion Coefficients: DJIA Portfolio

The table compares the average out-of-sample performance of the static and dynamic strategies for the DJIA portfolio using different risk aversion coefficients λ. A bootstrap procedure is applied to control for estimation error in expected returns. Panels A, B and C report results for the daily, weekly and monthly rebalanced portfolios, respectively. The static portfolios are constructed using the bootstrap unconditional means and covariance matrices. The table reports the annualised average returns (μ), the annualised average volatilities (σ), the average Sharpe ratios (SR), the p-values (proportions) that the dynamic strategies outperform the static alternatives in terms of Sharpe ratios, the average annualised performance fees Δ_γ (in basis points) that an investor with a constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolios, and the average breakeven transaction costs τ_γ (in basis points at the rebalancing frequency) that he is better off with the dynamic strategies.

	Static			*Long memory*			*Long memory vs. Static*					*OFLM3*			*OFLM3 vs. Static*				
λ	μ (%)	σ (%)	*SR*	μ (%)	σ (%)	*SR*	*p-value*	Δ_1	Δ_5	τ_1	τ_5	μ (%)	σ (%)	*SR*	*p-value*	Δ_1	Δ_5	τ_1	τ_5
Panel A. Daily rebalancing																			
1	4.62	3.20	0.20	5.49	5.08	0.29	0.73	79	48	2.1	1.3	5.19	2.84	0.41	0.91	59	64	5.6	6.2
2	4.31	1.61	0.20	4.76	2.55	0.29	0.77	43	35	2.3	1.9	4.59	1.42	0.41	0.92	28	29	5.3	5.6
3	4.20	1.06	0.20	4.51	1.69	0.30	0.77	30	27	2.4	2.1	4.39	0.94	0.41	0.91	19	19	5.4	5.6
4	4.15	0.79	0.20	4.38	1.27	0.30	0.75	22	20	2.4	2.2	4.29	0.71	0.41	0.91	14	14	5.4	5.5
5	4.12	0.64	0.20	4.31	1.01	0.30	0.77	18	17	2.4	2.2	4.24	0.56	0.42	0.92	11	12	5.5	5.6
Panel B. Weekly rebalancing																			
1	4.63	2.86	0.23	5.27	4.74	0.26	0.59	56	28	3.4	1.7	5.13	2.82	0.40	0.87	50	51	10.4	10.7
2	4.32	1.44	0.23	4.64	2.38	0.27	0.61	30	23	3.6	2.8	4.56	1.41	0.39	0.86	24	24	9.8	10.0
3	4.21	0.95	0.23	4.44	1.57	0.27	0.62	22	19	4.0	3.4	4.37	0.93	0.39	0.86	16	16	9.9	10.0
4	4.16	0.71	0.23	4.32	1.18	0.27	0.60	16	14	3.9	3.4	4.28	0.70	0.39	0.85	12	12	10.0	10.1
5	4.13	0.57	0.23	4.26	0.94	0.27	0.63	13	12	3.9	3.6	4.22	0.56	0.40	0.87	10	10	10.2	10.3
Panel C. Monthly rebalancing																			
1	4.63	2.81	0.23	5.14	3.77	0.30	0.66	48	35	9.6	7.2	5.17	2.65	0.44	0.90	55	57	35.4	37.5
2	4.32	1.42	0.23	4.58	1.89	0.31	0.68	26	23	10.3	9.1	4.58	1.32	0.44	0.91	27	27	34.1	35.3
3	4.21	0.93	0.23	4.39	1.25	0.31	0.69	18	16	10.7	9.9	4.38	0.87	0.43	0.91	18	18	34.5	35.2
4	4.16	0.70	0.23	4.29	0.95	0.30	0.67	13	12	10.5	9.9	4.29	0.66	0.43	0.91	13	14	34.7	35.2
5	4.13	0.56	0.23	4.23	0.75	0.31	0.70	11	10	10.6	10.1	4.23	0.53	0.44	0.91	11	11	35.5	35.9

Table 7.9. Comparison with Other Conditional Volatility Models: DJIA Portfolio

The table compares the out-of-sample performance of the OFLM3 portfolio with that of other benchmark portfolios constructed using different conditional volatility models for the DJIA dataset. A bootstrap procedure is applied to control for estimation error in expected returns. I generate an artificial sample of 4000 observations by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. The procedure is repeated with 1,000 trials. For each benchmark, the table reports the average annualised realised return (μ), the annualised realised volatility (σ), the average Sharpe ratio (SR), the *p*-value (proportion) out of 1000 trials that the OFLM3 portfolio outperforms the benchmark portfolio in terms of Sharpe ratios, the average abnormal return (*M2*) of the OFLM3 portfolio relative to the benchmark, and the average annualised performance fee Δ_γ (in basis points) that an investor with a constant relative risk coefficients of γ is willing to pay to switch from the benchmark to the factor volatility strategy. Turnover is the average value traded for all stocks in each period.

	***μ* (%)**	***σ* (%)**	***SR***	***p-value***	***M2* (%)**	Δ_1	Δ_5	***Turnover***
Panel A. Daily rebalancing								
OFLM3	5.193	2.840	0.414					0.044
<u>*Benchmarks*</u>								
EWMA	6.268	12.539	0.180	0.932	2.970	-31	274	0.508
GARCH-DCC	4.752	2.824	0.264	0.909	0.426	44	44	0.057
LM-EWMA-DCC	4.981	2.900	0.336	0.788	0.228	21	22	0.058
FIGARCH-DCC	4.295	8.474	0.051	0.953	3.214	124	263	0.705
CGARCH-DCC	4.802	2.948	0.269	0.889	0.426	39	41	0.071
Panel B. Weekly rebalancing								
OFLM3	5.132	2.822	0.396					0.097
<u>*Benchmarks*</u>								
EWMA	5.248	5.748	0.215	0.907	1.051	1	53	0.345
GARCH-DCC	4.702	2.766	0.252	0.900	0.403	43	42	0.102
LM-EWMA-DCC	4.772	2.969	0.260	0.890	0.417	36	38	0.114
FIGARCH-DCC	2.386	10.994	-0.129	0.992	6.015	338	591	0.634
CGARCH-DCC	4.788	2.846	0.276	0.875	0.345	34	35	0.100
Panel C. Monthly rebalancing								
OFLM3	5.171	2.648	0.439					0.130
<u>*Benchmarks*</u>								
EWMA	5.145	5.457	0.211	0.948	1.254	14	62	0.685
GARCH-DCC	4.750	2.453	0.306	0.862	0.332	42	39	0.136
LM-EWMA-DCC	4.967	2.516	0.387	0.723	0.140	20	19	0.140
FIGARCH-DCC	4.324	5.225	0.067	0.979	1.970	95	138	0.573
CGARCH-DCC	4.851	2.639	0.322	0.856	0.313	32	32	0.138

Table 7.10. Comparison with Other Conditional Volatility Models: International Stock and Bond Portfolio

The table compares the out-of-sample performance of the OFLM4 portfolio with that of other benchmark portfolios constructed using different conditional volatility models for the international stock and bond dataset. A bootstrap procedure is applied to control for estimation error in expected returns. I generate an artificial sample of 4000 observations by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. The procedure is repeated with 1,000 trials. For each benchmark, the table reports the average annualised realised return (μ), the annualised realised volatility (σ), the average Sharpe ratio (SR), the *p*-value (proportion) out of 1000 trials that the OFLM4 portfolio outperforms the benchmark portfolio in terms of Sharpe ratios, the average abnormal return (*M2*) of the OFLM4 portfolio relative to the benchmark, and the average annualised performance fee Δ_γ (in basis points) that an investor with a constant relative risk coefficients of γ is willing to pay to switch from the benchmark to the factor volatility strategy. Turnover is the average value traded for all stocks in each period.

	μ (%)	***σ (%)***	***SR***	***p-value***	***M2 (%)***	Δ_1	Δ_5	***Turnover***
Panel A. Weekly rebalancing								
OFLM4	8.096	2.670	1.537					0.23
Benchmarks								
EWMA	17.116	18.636	0.704	0.986	15.632	-715	40	4.98
GARCH-DCC	11.124	18.407	0.405	1.000	21.243	-117	628	14.90
LM-EWMA-DCC	11.715	16.207	0.489	1.000	17.280	-219	355	13.24
CGARCH-DCC	9.995	14.588	0.426	1.000	16.511	-75	388	7.49
Panel B. Monthly rebalancing								
OFLM4	7.924	2.502	1.573					0.45
Benchmarks								
EWMA	16.987	17.064	0.764	0.966	14.206	-733	-63	11.12
GARCH-DCC	10.063	14.954	0.434	1.000	17.758	-74	437	12.94
LM-EWMA-DCC	11.720	16.479	0.489	0.997	18.645	-205	450	17.43
CGARCH-DCC	9.510	12.156	0.505	0.999	13.795	-63	269	8.23

Table 7.11. Comparison of the Factor Models

The table compares the out-of-sample performance of portfolios constructed from different factor models. A bootstrap procedure is applied to control for estimation error in expected returns. I generate an artificial sample of 4000 observations by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. The procedure is repeated with 1,000 trials. The Factorα portfolio is constructed using the unconditional factor model with α factors. For each Factorα portfolio, the table reports the average annualised average return (μ), annualised average volatility (σ) and average Sharpe ratio (SR). For each pair of the Factor and OFLM portfolios, the table also reports the *p*-value (proportion) that the OFLMα portfolio outperforms the benchmark Factorα portfolio in the first column in terms of Sharpe ratios, the annualised average abnormal return (*M2*) of the OFLMα portfolio over the benchmark, the average performance fee Δ_γ (in basis points) that an investor with a constant relative risk coefficient of γ is willing to pay to switch from benchmark portfolio to the corresponding OFLM portfolio, and the average breakeven transaction cost τ_γ (in basis points at the rebalancing frequency) that he is better off with the dynamic OFLM strategy.

				Vs. OFLMα					
	μ (%)	***σ (%)***	***SR***	***p-value***	***M2 (%)***	Δ_1	Δ_5	τ_1	τ_5
Panel A. International stock and bond portfolio									
Weekly rebalancing									
Factor2	6.812	2.477	1.139	0.978	0.377	43.4	43.0	5.4	5.3
Factor3	7.145	2.596	1.215	0.942	0.320	34.1	34.0	3.9	3.9
Factor4	7.691	2.624	1.411	0.949	0.333	40.3	39.9	4.5	4.5
Monthly rebalancing									
Factor2	6.805	2.329	1.210	0.865	0.175	40.4	38.6	10.0	9.5
Factor3	7.144	2.403	1.314	0.774	0.147	33.0	31.6	7.4	7.1
Factor4	7.661	2.425	1.513	0.797	0.174	40.4	38.9	9.0	8.7
Panel B. DJIA portfolio									
Daily rebalancing									
Factor2	4.927	2.544	0.357	0.355	-0.127	-6.9	-9.3	–	–
Factor3	5.084	2.666	0.398	0.556	0.045	10.5	8.6	1.1	0.9
Factor4	5.041	2.688	0.381	0.643	0.146	26.8	23.4	2.7	2.3
Weekly rebalancing									
Factor2	4.918	2.467	0.366	0.246	-0.217	-14.0	-17.4	–	–
Factor3	5.058	2.572	0.405	0.481	-0.020	6.7	4.0	1.4	0.8
Factor4	4.996	2.585	0.381	0.524	0.024	14.9	11.1	3.1	2.3
Monthly rebalancing									
Factor2	4.918	2.303	0.394	0.290	-0.146	-5.6	-8.9	–	–
Factor3	5.055	2.387	0.435	0.512	0.008	11.0	8.3	7.8	5.8
Factor4	5.044	2.470	0.418	0.563	0.061	19.8	16.2	13.8	11.3

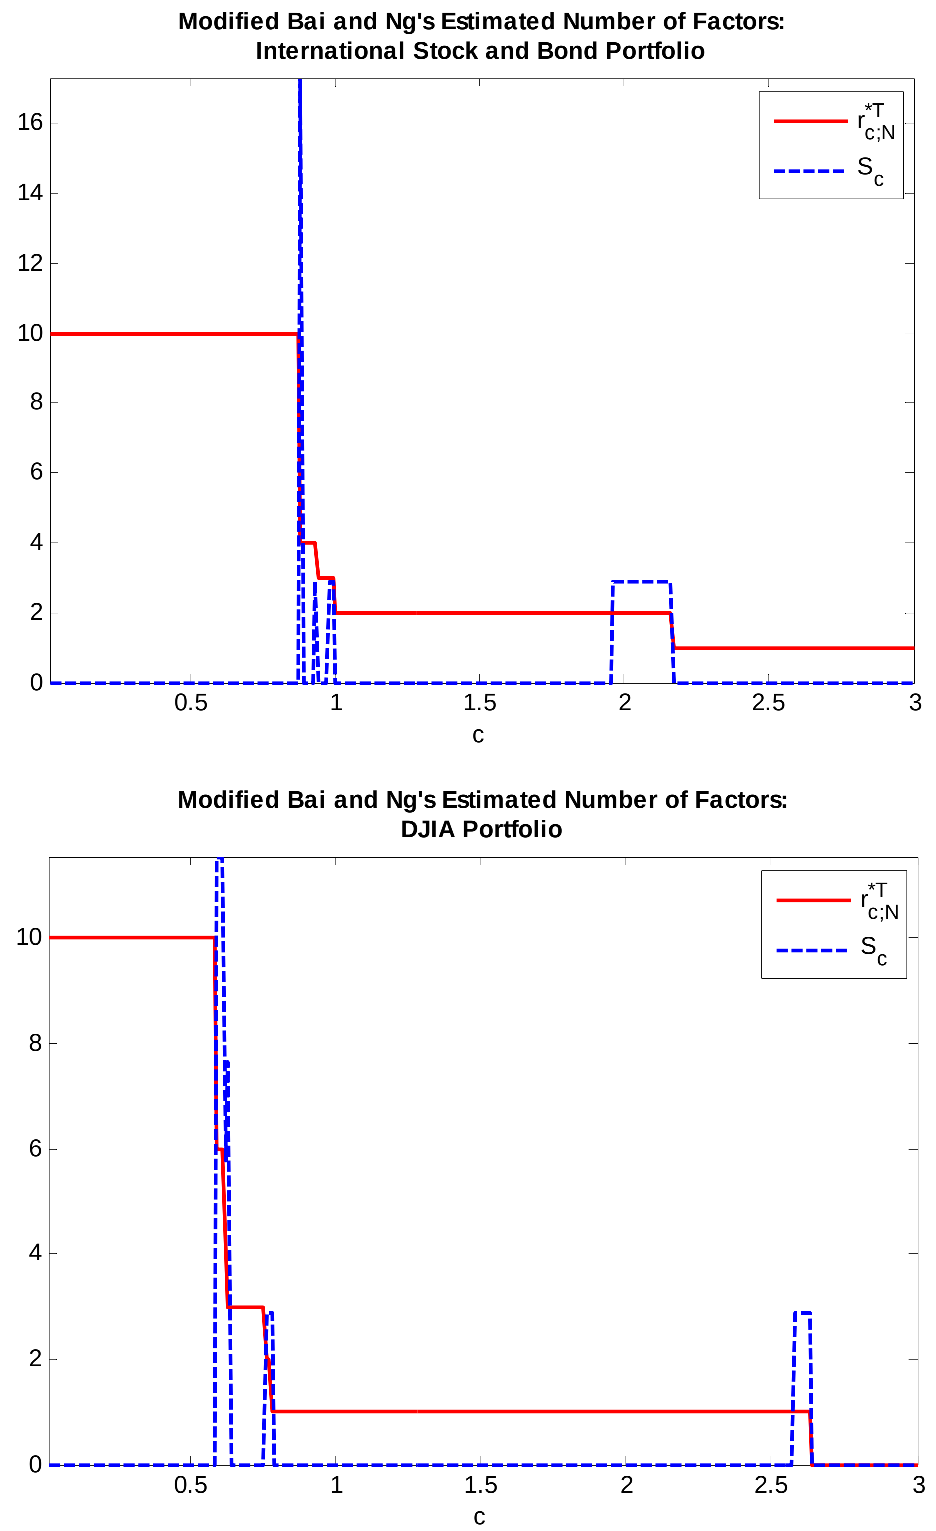

Figure 7.1. Determining the Number of Common Factors.

The number of factors is determined using the modified Bai and Ng's (2002) information criterion of Alessi et al. (2007). The number of factors corresponds to the second stable region, i.e., the plateau of the solid line associated with the second flat zero-level dashed line.

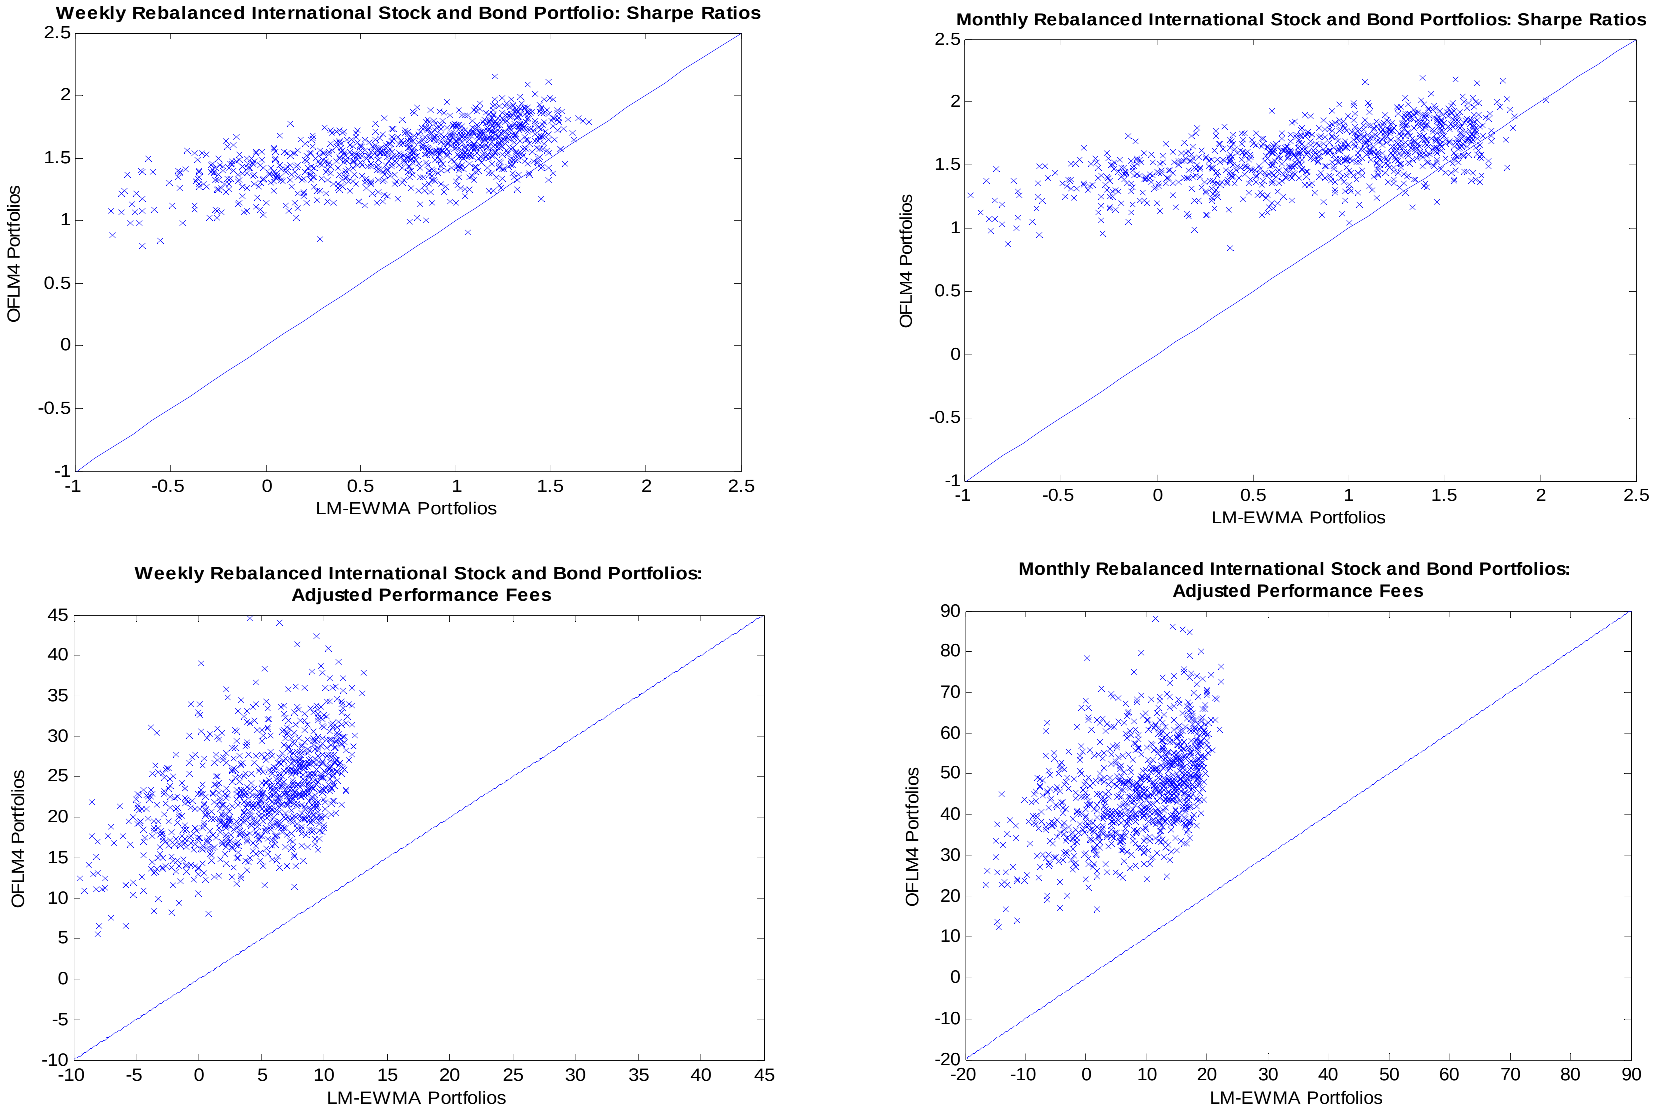

Figure 7.2. The Sharpe Ratios and Adjusted Performance Fees of the Bootstrap LM-EWMA and OFLM4 Portfolios

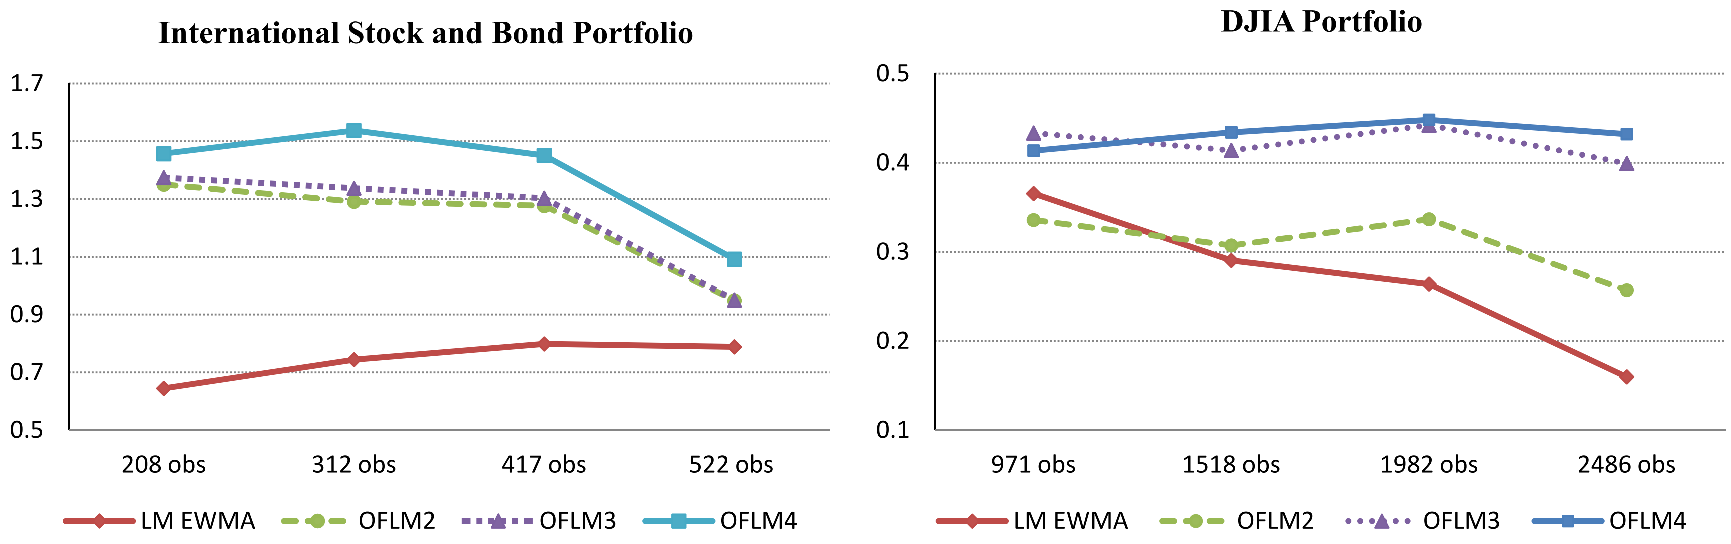

Figure 7.3. Sensitivity to Estimation Window: The Sharpe Ratios of the Dynamic Portfolios.

The figure plots the average Sharpe ratios of the optimal portfolios constructed from the LM-EWMA and OFLM models with different estimation windows. A wide range of bootstrap vectors of expected returns is employed to account for estimation error in expected returns. The estimation windows correspond to 4, 6, 8 and 10 years of weekly and daily data.

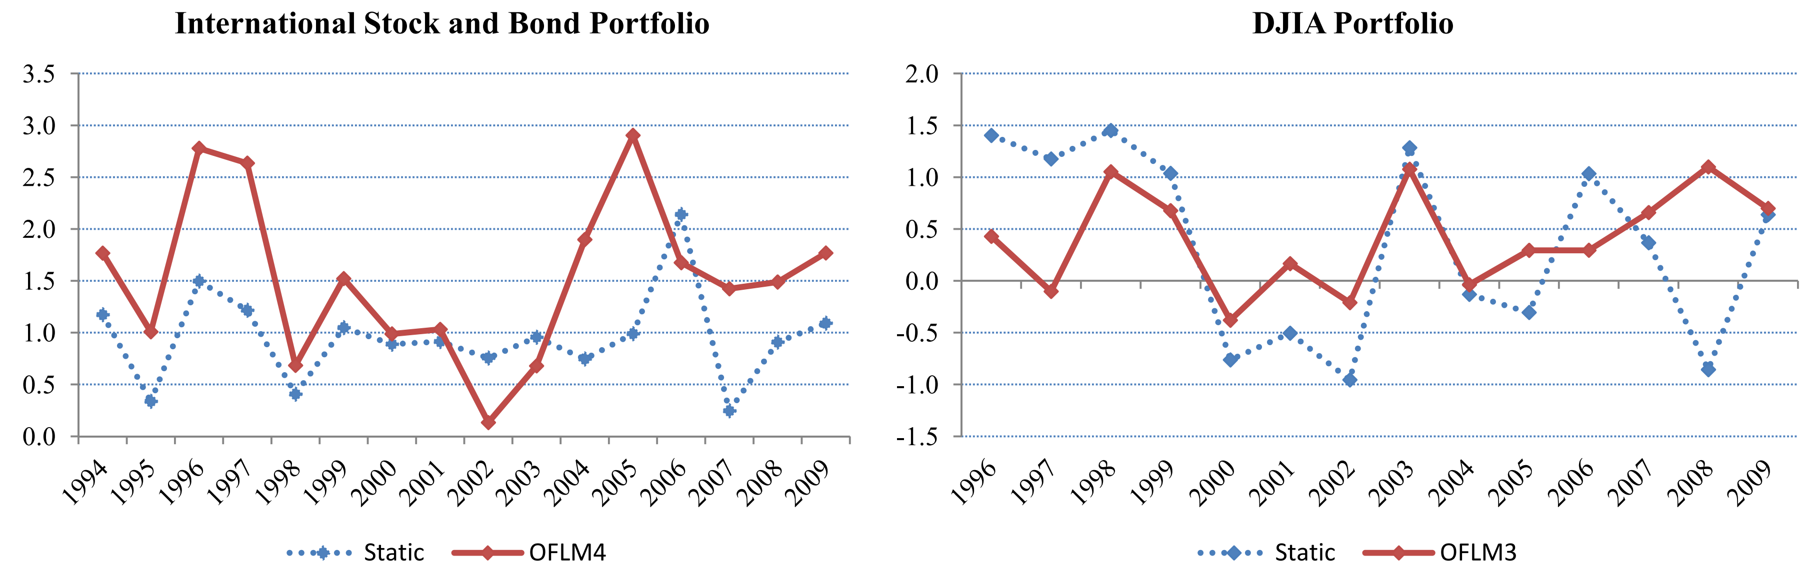

Figure 7.4. Average Sharpe Ratios of the Static and Dynamic Factor Long Memory Portfolios over Years.

The figure plots the average Sharpe ratios of the static and the dynamic factor long memory portfolios over years. A wide range of bootstrap vectors of expected returns is employed to account for estimation error in expected returns.

Appendix 7.1. Comparison of the Orthogonal Factor Long Memory EWMA and the Orthogonal Factor EWMA Models

Similar to the OFLM model, the Orthogonal Factor EWMA, the OF-EWMA model is developed by embedding the short memory EWMA model in an orthogonal factor structure. Unlike the Orthogonal EWMA model of Alexander (2001), the Orthogonal Factor EWMA model assumes both the volatilities of the factors and the residuals follow EWMA processes. The appendix compares the out-of-sample performance of the OF-EWMA and OFLM portfolios to evaluate the benefits of allowing for long memory vs. short memory volatility in the factor structure.

A bootstrap procedure is applied to control for estimation error in expected returns. I generate an artificial sample of 4000 observations by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. The procedure is repeated with 1,000 trials. For each OF-EWMA portfolio, the table below reports the average annualised average return (μ), annualised average volatility (σ) and Sharpe ratio (SR). For each pair of the OF-EWMAα and OFLMα portfolios, where α is the number of factors, the table also reports the *p*-value (proportion) that the OFLMα portfolio outperforms the benchmark OF-EWMAα portfolio in the first column in terms of Sharpe ratio, the annualised average abnormal return *M2* of the OFLMα portfolio over the benchmark, and the average performance fee Δ_γ (in basis points) that an investor with the constant relative risk coefficient of γ is willing to pay to switch from benchmark portfolio to the corresponding OFLM portfolio. The table also reports the average turnovers of the OF-EWMA and the OFLM strategies.

The results suggest the outperformance of the long memory factor OFLM model relative to the short memory factor OF-EWMA model, especially in the DJIA portfolio. With the DJIA dataset, the OFLM model consistently produces portfolios with higher Sharpe ratios than those produced by the OF-EWMA models in more than 85% of the trials. The OFLM model performs less remarkably in the international stock and bond portfolio. However, the OFLM4 portfolio still dominates the OF-EWMA4 portfolio in terms of the Sharpe ratio.

	μ (%)	σ (%)	SR	*Vs. OFLMα* p-value	M2 (%)	Δ_1	Δ_5	*Turnover* OF EWMA	OFLM
Panel A. International stock and bond portfolio									
Weekly rebalancing									
OF-EWMA2	7.354	2.618	1.284	0.615	0.018	-10	-9	9.9	20.0
OF-EWMA3	7.630	2.703	1.346	0.388	-0.021	-14	-13	11.0	22.0
OF-EWMA4	8.204	2.765	1.524	0.695	0.038	-11	-9	12.1	23.1
Monthly rebalancing									
OF-EWMA2	7.378	2.626	1.295	0.307	-0.023	-16	-15	22.6	38.5
OF-EWMA3	7.665	2.654	1.387	0.296	-0.034	-18	-17	25.4	42.9
OF-EWMA4	8.224	2.698	1.572	0.685	0.036	-15	-14	27.1	44.7
Panel B. DJIA portfolio									
Daily rebalancing									
OF-EWMA2	4.859	3.161	0.268	0.884	0.125	2	6	4.3	4.2
OF-EWMA3	5.247	3.206	0.383	0.867	0.099	-4	0	4.5	4.4
OF-EWMA4	5.358	3.346	0.399	0.879	0.118	-3	2	4.8	4.7
Weekly rebalancing									
OF-EWMA2	4.700	2.935	0.237	0.941	0.130	9	11	5.9	9.2
OFEWMA3	5.051	2.977	0.349	0.945	0.143	9	10	6.2	9.7
OFEWMA4	5.091	3.076	0.349	0.905	0.123	7	9	6.6	10.3
Monthly rebalancing									
OF-EWMA2	4.822	2.829	0.291	0.861	0.118	5	8	13.5	12.4
OF-EWMA3	5.150	2.860	0.399	0.862	0.112	3	5	14.0	13.0
OF-EWMA4	5.263	3.010	0.415	0.776	0.082	-1	2	14.9	13.9

Appendix 7.2. Comparison of the Orthogonal Factor Long Memory EWMA and the Orthogonal Factor GARCH Models

Similar to the OF-EWMA model described above, the Orthogonal Factor GARCH, the OF-GARCH model is achieved by embedding the short memory GARCH model in an orthogonal factor structure. Unlike the Orthogonal GARCH model of Alexander (2001), the Orthogonal Factor GARCH model assumes both the volatilities of the factors and of the residuals follow GARCH processes. This appendix uses OF-GARCH model as another benchmark to evaluate the benefits of allowing for long memory vs. short memory volatility in the factor structure.

A bootstrap procedure is applied to control for estimation error in expected returns. I generate an artificial sample of 4000 observations by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. The procedure is repeated with 1,000 trials. Again, for each OF-GARCH portfolio, the table below reports the average annualised average return (μ), annualised average volatility (σ) and Sharpe ratio (SR). For each pair of the OF-GARCHα and OFLMα portfolios, where α is the number of factors, the table also reports the *p*-value (proportion) that the OFLMα portfolio outperforms the benchmark OF-GARCHα portfolio in the first column in terms of Sharpe ratio, the annualised average abnormal return *M2* of the OFLMα portfolio over the benchmark, and the average performance fee Δ_γ (in basis points) that an investor with the constant relative risk coefficient of γ is willing to pay to switch from benchmark portfolio to the corresponding OFLM portfolio. The table also reports the average turnovers of the OF-EWMA and the OFLM strategies.

While the long memory factor OFLM model produces better portfolios in terms of Sharpe ratios and performance fees than the short memory factor OF-GARCH model in the international stock and bond portfolio, it underperforms in the DJIA portfolio. With the DJIA dataset, though the OFLM portfolios with 3 and 4 factors still generate higher performance fees, they have lower Sharpe ratios and slightly higher turnover.

	μ (%)	*σ (%)*	*SR*	*Vs. OFLMα*				*Turnover*	
				p-value	*M2 (%)*	Δ_1	Δ_5	*OF EWMA*	*OFLM*
Panel A. International stock and bond portfolio									
Weekly rebalancing									
OF-GARCH2	7.071	2.482	1.241	0.855	0.126	18	17	26.7	20.0
OF-GARCH3	7.335	2.622	1.274	0.908	0.166	15	15	28.7	22.0
OF-GARCH4	7.951	2.645	1.497	0.810	0.107	14	14	30.5	23.1
Monthly rebalancing									
OF-GARCH2	7.036	2.423	1.260	0.695	0.062	18	17	32.9	38.5
OF-GARCH3	7.303	2.518	1.317	0.894	0.144	17	17	36.2	42.9
OF-GARCH4	7.924	2.502	1.573	0.604	0.029	14	14	37.8	44.7
Panel B. DJIA portfolio									
Daily rebalancing									
OF-GARCH2	4.874	2.581	0.335	0.323	-0.070	-2	-4	4.3	4.2
OF-GARCH3	5.117	2.660	0.415	0.487	0.001	7	5	4.3	4.4
OF-GARCH4	5.242	2.785	0.439	0.465	-0.012	7	5	4.4	4.7
Weekly rebalancing									
OF-GARCH2	4.815	2.559	0.317	0.248	-0.093	-3	-6	7.4	9.2
OF-GARCH3	5.051	2.625	0.396	0.495	0.002	8	5	7.5	9.7
OF-GARCH4	5.117	2.718	0.406	0.367	-0.045	3	1	7.8	10.3
Monthly rebalancing									
OF-GARCH2	4.890	2.315	0.385	0.200	-0.121	-3	-6	9.6	12.4
OF-GARCH3	5.092	2.392	0.453	0.393	-0.034	7	5	9.9	13.0
OF-GARCH4	5.222	2.543	0.475	0.223	-0.082	2	-1	10.7	13.9

Chapter 8

Conclusions

8.1 Conclusions

It is well established that the covariance matrix of short horizon financial asset returns is both time varying and highly persistent. A number of multivariate conditional volatility models, including the multivariate RiskMetrics EWMA model and multivariate GARCH models have been developed to capture these features. Conditional volatility models have been found to produce impressive estimates and forecasts of the covariance matrix, which are now routinely used in many areas of applied finance, including asset allocation. The literature suggests that dynamic asset allocation strategy benefits from exploiting the forecasts of multivariate conditional volatility models relative to using the unconditional covariance matrix (see, for example, Fleming et al., 2001, Fleming et al., 2003, Engle and Colacito, 2006, Thorp and Milunovich, 2007). Recent evidence also suggests that volatility has longer memory than that implied in the EWMA and GARCH models (see, for example, Taylor, 1986, Ding et al., 1993, Andersen et al., 2001) and long memory volatility models generally outperform short memory volatility models in terms of forecast performance in both univariate and multivariate context (Vilasuso, 2002, Zumbach, 2006, Niguez and Rubia, 2006). Multivariate long memory volatility models, with slow decaying autocorrelations, are designed to capture the high persistence feature of volatility and covariance and exploit this feature to provide more reasonable forecasts of the covariance matrix over long horizons. Allowing for long memory volatility dynamics in forecasts of the covariance matrix, therefore, may bring potential benefits for asset allocation. However, owing to the complexity in estimation of long memory volatility models, the literature on multivariate long memory volatility modelling has restricted itself to the analysis of low dimensional covariance matrices, and has provided limited evidence on the benefits from allowing for long memory dynamics in the multivariate setting. So far multivariate long memory conditional volatility models have not been used in the asset allocation framework, where forecasts of the high dimensional covariance matrix over long horizons are normally required.

The thesis undertakes a detailed analysis of multivariate long memory conditional volatility models and their application in dynamic asset allocation. The analysis is conducted in a systematic approach, starting from the evaluation of the forecast performance of long memory conditional volatility models, to the study of their benefits in the dynamic asset allocation framework, and finally to the use of a factor structure to reduce estimation error in estimating the high dimensional long memory covariance matrix. While many alternative volatility models have been developed in the literature, the thesis purposely chooses to employ parsimonious models that can be used to forecast high dimensional covariance matrices. In particular, the four multivariate long memory volatility models (the long memory EWMA, long memory EWMA-DCC, FIGARCH-DCC and Component GARCH-DCC models) are evaluated against one another and against the two multivariate short memory volatility models (the short memory EWMA and GARCH-DCC models). The thesis employs four datasets, both low and high dimensional covariance matrices with both low and high correlation assets and in both the US and the international markets. The analysis is original and distinctive as this is the first study to evaluate a wide range of multivariate long memory volatility models in high dimensional systems for asset allocation.

The findings consistently show that it is beneficial to allow for long memory volatility dynamics in estimating the covariance matrix of returns for asset allocation over both short and long horizons. First, multivariate long memory volatility models generally produce forecasts of the covariance matrix that are statistically more accurate and informative, and economically more useful than those produced by short memory volatility models. Second, investors are better off constructing portfolios with the long memory volatility models than with the static unconditional and the short memory volatility alternatives. Third, embedding a factor structure in the long memory volatility framework to reduce estimation error in forecasts of the covariance matrix brings substantial benefits. The dynamic factor long memory volatility timing strategy systematically dominates the static and other dynamic volatility timing strategies that employ both short memory and long memory volatility models. The long memory factor model also outperforms the traditional unconditional factor and the short memory factor models. More details of the findings are given as follows.

First, the forecast performance of multivariate long memory conditional volatility models are evaluated among themselves and against that of short memory conditional volatility models. Popular statistical measures are applied to measure the accuracy, bias and information content of forecasts of alternative volatility models. Given the well-documented drawbacks of the statistical criteria, the analysis also employs forecast performance measures based on economic loss criteria. Specifically, the research studies the usefulness of forecasts of the conditional covariance matrix in the asset allocation framework of Engle and Colacito (2006). The results consistently show that multivariate long memory conditional volatility models generally produce better forecasts of the covariance matrix, in terms of both statistical and economic measures than short memory volatility models. Also, among the long memory volatility models, the two parsimonious long memory EWMA models dominate in a majority of cases across all forecast horizons of up to 3 months. Although the Component GARCH and FIGARCH models are also designed to capture long memory volatility, their performance is less impressive, which may be attributable to their high degree of parameterisation and complex estimation procedure. The results are robust to different investment horizons and estimation windows. The findings of the analysis are consistent with those in the univariate volatility literature.

Long memory conditional covariance matrices have been found to produce optimal portfolios with lower realised volatilities than static unconditional and short memory covariance matrices. However, it does not follow that the portfolio with the lowest volatility is necessarily the best portfolio in terms of portfolio performance measures such as the Sharpe ratio. Consequently, the research continues to further explore the value of long memory conditional volatility models, studying the economic benefits accruing to volatility timing strategies using the framework of Fleming et al. (2001). Portfolios constructed from the six multivariate conditional volatility models, both long memory and short memory, are evaluated using popular performance measures such as the out-of-sample Sharpe ratio, the abnormal return and the performance fee. The research consistently identifies the gains of incorporating long memory volatility dynamics in forecasts of the covariance matrix for asset allocation. The long memory conditional volatility models, especially the parsimonious long memory EWMA models, generally produce better portfolios than the static unconditional and the short memory volatility models at all investment horizons. The research also studies the effects of transaction costs in conducting the dynamic volatility timing strategies. With the presence of transaction costs, the gains from the daily rebalanced portfolios deteriorate, however, it is still worth implementing the dynamic strategies at lower rebalancing frequencies. Again, the two long memory EWMA models are consistently the most favourable models in terms of low transaction costs. The results apply to the four low and high dimension datasets and are robust to estimation error in expected returns, the choice of risk aversion coefficient and the use of a long-only constraint.

Finally, the research deals with the estimation error inherent in estimating the covariance matrix. Estimation error in the long memory volatility covariance matrix may be more excessive than in the unconditional covariance matrix, not least because of the high degree of parameterisation. The estimation error problem gets even worse when the high dimensional covariance matrix typically encountered in asset allocation are inversed. The thesis, therefore, imposes a factor structure to the long memory volatility framework to control for estimation error in estimating the high dimensional covariance matrix. In so doing, the research develops a dynamic factor long memory conditional volatility, the Orthogonal Factor Long Memory or OFLM, model by combining the univariate long memory EWMA model of Zumbach (2006) with an orthogonal factor structure. The new factor model follows richer processes than normally assumed, in which both the factors and idiosyncratic shocks are modelled with long memory behaviour in their volatilities. The factor-structured OFLM model is evaluated against the six other multivariate conditional volatility models, especially the fully estimated multivariate long memory EWMA model of Zumbach (2009b), in terms of forecast performance and economic benefits. The results suggest that the OFLM model generally produces impressive forecasts over both short and long forecast horizons. In the volatility timing framework, portfolios constructed with the OFLM model consistently dominate the static and other dynamic volatility timing portfolios at all rebalancing frequencies. In particular, the outperformance of the factor OFLM model to the fully estimated LM-EWMA model evidently confirms the advantage of the factor structure in reducing estimation error. Employing the factor structure also significantly reduces transaction costs, making dynamic trading more feasible. The findings also suggest that the long memory factor model generally produces better portfolios than the traditional unconditional factor and the short memory factor models. The results are, again, robust to estimation error in expected returns, the choice of risk aversion coefficient, and the length of estimation window.

8.2 Limitations of the Research

To evaluate the forecast performance of the long memory conditional volatility models, the research employs both statistical and economic criteria. The statistical measures such as the RMSE, MAE or Mincer- Zarnowitz regression use the squares and cross-products of daily returns as proxies for the actual variances and covariances being forecast. The use of low frequency daily volatility as a proxy for true volatility introduces considerable noise that inflates the forecast errors of the conditional volatility forecasts, substantially reducing their explanatory power. Given this problem with the daily volatility proxy, it would be more appropriate if the research could employ other higher frequency volatilities, e.g., realised volatility or range-based volatility, as proxies for the true volatility. Also, it would be interesting to compare the economic forecast performance of the long memory volatility models with that of the recently popular high frequency volatility models.

Though the analysis covers four datasets, both low and high dimensional portfolios with both low and high correlation assets, the assets consists of only stocks and bonds. It may be of interest to study the performance of long memory volatility models with other assets, e.g., foreign exchange rates or commodity futures.

The economic value of the multivariate conditional volatility models are evaluated in a rather restrictive volatility timing framework, in which expected returns are assumed constant and investors myopically update their portfolios based on forecasts of the covariance matrix in every period. This may not correspond to the real-world practice where expected returns are also time varying and investors may be more concerned with long term investments. Besides, apart from asset allocation, the use of multivariate conditional volatility models may bring potential benefits to other practical problems, e.g., risk management. The analysis of conditional volatility models, therefore, can be extended to other contexts.

The dynamic factor model assumes constant factor loadings over time. However, betas are affected by the covariances between the factors and asset returns and the volatilities of the factors, which are both time-varying. It would thus be of interest to relax the constant beta assumption, developing a dynamic factor model with time-varying conditional betas. Also, expected returns may be time-varying and dynamic autoregressive factor models may be developed to estimate both expected returns and covariance matrix.

8.3 Suggestions for Future Research

The methodology and findings of the thesis suggest some implications for future research. First, the non-DCC conditional covariance matrix estimators (such as the EWMA model with exponential weights and the LM-EWMA model with logarithmic weights) impose the same dynamic structure on all elements of the covariance matrix, which facilitates their implementation in high dimensional systems, but it comes at a cost in terms of estimation error. A potentially less correctly specified but more flexible DCC structure is generally found to perform better in high dimension and high correlation systems. It would be interesting to investigate this issue in greater detail.

Second, the long memory DCC framework models the dynamic processes of volatility and correlation separately, using long memory volatility models for individual volatility. However, the DCC models assume a short memory process for correlation. A more appropriate approach would be to apply long memory dynamics for both volatility and correlation, embedding the long memory volatility models in a long memory correlation framework such as, e.g., the DCC-MIDAS model of Colacito et al. (2011).

Third, the analysis of the benefits of allowing for conditional covariance matrix estimators in asset allocation can be extended in several ways. The assumption of constant expected returns can be relaxed. Time-varying volatility affects returns and it is hence not justifiable to separate the movement of expected returns with those of volatility and correlation. However, research then has to differentiate between the effects of better estimates of expected returns and better estimates of the covariance matrix in the improvement of the optimal portfolios. Also, it may be useful to study the economic value of dynamic strategies in an intertemporal asset allocation framework. Dynamic strategies may behave differently in the presence of hedging demands.

Fourth, estimates of the covariance matrix are inherently subject to estimation error. Moreover, some of the eigenvalues of the high dimensional covariance matrix are inevitably very small, and so the inverse of the covariance matrix used in the asset allocation is likely to be ill-conditioned. The research employs a statistical factor structure to reduce estimation error in estimating the long memory conditional covariance matrix. Long memory volatility models can also be embedded in other macroeconomic and fundamental factor models with time-varying betas. Besides, other structures, such as the Bayesian shrinkage models, can also be applied to produce a robust estimate of the long memory conditional covariance matrix.

Finally, the analysis of the limitations of the thesis also reveals some other promising directions for further research. They include, for example, the use of high frequency volatility models as the benchmarks, the coverage of other assets and the extension of the analysis to other practical problems.

References

ALESSI, L., BARIGOZZI, M. & CAPASSO, M. 2007. A robust criterion for determining the number of static factors in approximate factor models. *LEM Working Papers 2007/19, LEM (Laboratory of Economics and Management), Sant's Anna School of Advanced Studies, Pisa*.

ALEXANDER, C. 2001. Orthogonal Garch. *Mastering Risk,* 2**,** 21--38.

ANDERSEN, T. G., BOLLERSLEV, T., CHRISTOFFERSEN, P. F. & DIEBOLD, F. X. 2006. Volatility and correlation forecasting. *In:* ELLIOTT, G., GRANGER, C. & TIMMERMANN, A. (eds.) *{Handbook of Economic Forecasting}.* North Holland.

ANDERSEN, T. G., BOLLERSLEV, T., DIEBOLD, F. X. & EBENS, H. 2001. The distribution of realized stock return volatility. *The Journal of Financial Economics,* 61**,** 43--76.

ANEJA, Y. P., CHANDRA, R. & GUNAY, E. 1989. A portfolio approach to estimating the average correlation coefficient for the constant correlation model. *The Journal of Finance,* 44**,** 1435--1438.

ANG, A., HODRICK, R. J., XING, Y. & ZHANG, X. 2006. The Cross-Section of Volatility and Expected Returns. *The Journal of Finance,* 61**,** 259--299.

ANG, A., HODRICK, R. J., XING, Y. & ZHANG, X. 2009. High idiosyncratic volatility and low returns: International and further US evidence. *The Journal of Financial Economics,* 91**,** 1--23.

ANGERER, X. & LAM, P. S. 2009. Income Risk and Portfolio Choice: An Empirical Study. *The Journal of Finance,* 64**,** 1037--1055.

BAI, J. & NG, S. 2002. Determining the number of factors in approximate factor models. *Econometrica,* 70**,** 191--221.

BAILLIE, R. T. 1996. Long Memory Processes and Fractional Integration in Econometrics. *Journal of Econometrics,* 73**,** 5-59.

BAILLIE, R. T., BOLLERSLEV, T. & MIKKELSEN, H. O. 1996. Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics,* 74**,** 3--30.

BARBERIS, N. 2000. Investing for the long run when returns are predictable. *The Journal of Finance,* 55**,** 225--264.

BARRY, C. B. 1974. Portfolio analysis under uncertain means, variances, and covariances. *The Journal of Finance*, 29**,** 515--522.

BENGTSSON, C. & HOLST, J. 2002. On portfolio selection: Improved covariance matrix estimation for Swedish asset returns. *Working Paper*.

BEST, M. J. & GRAUER, R. R. 1991. On the sensitivity of mean-variance-efficient portfolios to changes in asset means: some analytical and computational results. *The Review of Financial Studies,* 4**,** 315--342.

BLACK, F. 1976. Studies of stock price volatility changes. *Proceeding of the Meeting of Business and Economic Statistics Section.*

BLACK, F. & LITTERMAN, R. 1992. Global portfolio optimization. *Financial Analysts Journal,* 48**,** 28--43.

BOLLERSLEV, T. 1986. Generalised autoregressive conditional heteroskedasticity. *Journal of Econometrics,* 31**,** 307--327.

BOLLERSLEV, T. 1990. Modelling the Coherence in Short-run Nominal Exchange Rates: A Multivariate Generalized ARCH Model. *The Review of Economics and Statistics,* 72**,** 498--505.

BOLLERSLEV, T. & GHYSELS, E. 1996. Periodic Autoregressive Conditional Heteroscedasticity. *Journal of Business & Economic Statistics,* 14**,** 139--151.

BRANDT, M. W. 2009. Portfolio choice problems. *In:* AIT-SAHALIA, Y. & HANSEN, L. P. (eds.) *Handbook of Financial Econometrics.* North Holland.

BRANDT, M. W., GOYAL, A., SANTA-CLARA, P. & STROUD, J. R. 2005. A simulation approach to dynamic portfolio choice with an application to learning about return predictability. *The Review of Financial Studies,* 18**,** 831--873.

BREIDT, F. J., CRATO, N. & DE LIMA, P. 1998. The detection and estimation of long memory in stochastic volatility. *Journal of Econometrics,* 83**,** 325--348.

BRINER, B. G. & CONNOR, G. 2008. How much structure is best? A comparison of market model, factor model and unstructured equity covariance matrices. *Journal of Risk,* 10**,** 3--30.

BRITTEN-JONES, M. 1999. The sampling error in estimates of mean-variance efficient portfolio weights. *The Journal of Finance,* 54**,** 655--671.

BROADIE, M. 1993. Computing efficient frontiers using estimated parameters. *Annals of Operations Research,* 45**,** 21--58.

BROWN, S. 1979. The effect of estimation risk on capital market equilibrium. *The Journal of Financial and Quantitative Analysis,* 14**,** 215--220.

BURMEISTER, E., ROLL, R. & ROSS, S. A. 2003. Using macroeconomic factors to control portfolio risk. *BIRR Portfolio Analysis, Inc.*

CAMPBELL, J. Y., COCCO, J., GOMES, F., MAENHOUT, P. J. & VICEIRA, L. M. 2001. Stock market mean reversion and the optimal equity allocation of a long-lived investor. *Review of Finance,* 5**,** 269--293.

CAMPBELL, J. Y. & VICEIRA, L. M. 2002. *Strategic asset allocation: portfolio choice for long-term investors*, Oxford University Press, USA.

CAPPIELLO, L., ENGLE, R. F. & SHEPPARD, K. K. 2003. Asymmetric Dynamics in the Correlations of Global Equity and Bond Returns. *SSRN eLibrary*.

CARHART, M. M. 1997. On persistence in mutual fund performance. *The Journal of Finance,* 52**,** 57--82.

CHAMBERLAIN, G. 1983. Funds, Factors, and Diversification in Arbitrage Pricing Models. *Econometrica,* 51**,** 1305--1323.

CHAN, L., KARCESKI, J. & LAKONISHOK, J. 1999. On portfolio optimization: Forecasting covariances and choosing the risk model. *The Review of Financial*

CHEN, N. F., ROLL, R. & ROSS, S. A. 1986. Economic forces and the stock market. *Journal of Business,* 59**,** 383--403.

CHUNHACHINDA, P., DANDAPANI, K., HAMID, S. & PRAKASH, A. J. 1997. Portfolio selection and skewness: Evidence from international stock markets. *The Journal of Banking & Finance,* 21**,** 143--167.

COCCO, J. F., GOMES, F. J. & MAENHOUT, P. J. 2005. Consumption and portfolio choice over the life cycle. *The Review of Financial Studies,* 18**,** 491.

COLACITO, R., ENGLE, R. F. & GHYSELS, E. 2011. A component model for dynamic correlations. *Journal of Econometrics,* 164, 45-59

CONNOR, G. 1995. The three types of factor models: A comparison of their explanatory power. *Financial Analysts Journal,* 51**,** 42--46.

CONNOR, G. & KORAJCZYK, R. A. 1993. A test for the number of factors in an approximate factor model. *The Journal of Finance,* 48**,** 1263--1291.

DAMMON, R. M., SPATT, C. S. & ZHANG, H. H. 2001. Optimal consumption and investment with capital gains taxes. *The Review of Financial Studies,* 14**,** 583--616.

DAVIDSON, J. 2004. Moment and memory properties of linear conditional heteroscedasticity models, and a new model. *Journal of Business & Economic Statistics,* 22**,** 16--29.

DAVIS, M. & NORMAN, A. 1990. Portfolio selection with transaction costs. *Mathematics of Operations Research,* 15**,** 676--713.

DEMIGUEL, V., GARLAPPI, L. & UPPAL, R. 2009. Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy? *The Review of Financial Studies,* 22**,** 1915--1953.

DEMIGUEL, V. & UPPAL, R. 2005. Portfolio investment with the exact tax basis via nonlinear programming. *Management Science,* 51**,** 277--290.

DIEBOLD, F. X. & MARIANO, R. S. 1995. Comparing Predictive Accuracy. *Journal of Business & Economic Statistics,* 13**,** 253--263.

DING, Z., GRANGER, C. W. J. & ENGLE, R. F. 1993. A long memory property of stock market return and a new model. *Journal of Empirical Finance,* 1**,** 83--106.

DISATNIK, D. J. & BENNINGA, S. 2007. Shrinking the covariance matrix. *The Journal of Portfolio Management,* 33**,** 55--63.

ELTON, E. J. & GRUBER, M. J. 1973. Estimating the Dependence Structure of Share Prices--Implications for Portfolio Selection. *The Journal of Finance***,** 1203--1232.

ENGLE, R. F. 1982. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica,* 50**,** 987--1007.

ENGLE, R. F. 2002. Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics,* 20**,** 339--350.

ENGLE, R. F. 2009. High dimension dynamic correlations. *In:* HENDRY, D. F., CASTLE, J. & SHEPHARD, N. (eds.) *The Methodology and Practice of Econometrics: A Festschrift in Honour of David F. Hendry.*

ENGLE, R. F. & BOLLERSLEV, T. 1986. Modelling the persistence of conditional variances. *Econometric Reviews,* 5**,** 1--50.

ENGLE, R. F. & COLACITO, R. 2006. Testing and valuing dynamic correlations for asset allocation. *Journal of Business and Economic Statistics,* 24**,** 238--253.

ENGLE, R. F. & LEE, G. 1999. A Permanent and Transitory Component Model of Stock Return Volatility *In:* ENGLE, R. & WHITE, H. (eds.) *Cointegration, Causality, and Forecasting: A Festschrift in Honor of Clive WJ Granger.* Oxford University Press.

ENGLE, R. F., LILIEN, D. M. & ROBINS, R. P. 1987. Estimating time varying risk premia in the term structure: The ARCH-M model. *Econometrica: Journal of the Econometric Society,* 55**,** 391--407.

ENGLE, R. F., KANE, A. & NOH, J. 1996. Index-option pricing with stochastic volatility and the value of accurate variance forecasts. *Review of Derivatives Research,* 1**,** 139--157.

ENGLE, R. F. & KRONER, K. F. 1995. Multivariate simultaneous generalized ARCH. *Econometric Theory,* 11**,** 122--150.

ENGLE, R. F. & NG, V. K. 1993. Measuring and testing the impact of news on volatility. *The Journal of Finance,* 48**,** 1749--1778.

ENGLE, R. F. & SHEPPARD, K. 2001. Theoretical and empirical properties of dynamic conditional correlation multivariate GARCH. *NBER Working Paper*.

FAMA, E. F. 1965. The Behavior of Stock Market Prices. *The Journal of Business,* 38**,** 34--105.

FAMA, E. F. & FRENCH, K. R. 1992. The cross-section of expected stock returns. *The Journal of Finance,* 47**,** 427--465.

FAMA, E. F. & FRENCH, K. R. 1993. Common risk factors in the returns on stocks and bonds. *The Journal of Financial Economics,* 33**,** 3--56.

FAN, J., FAN, Y. & LV, J. 2008. High dimensional covariance matrix estimation using a factor model. *Journal of Econometrics,* 147**,** 186--197.

FLEMING, J., KIRBY, C. & OSTDIEK, B. 2001. The economic value of volatility timing. *The Journal of Finance,* 56**,** 329--352.

FLEMING, J., KIRBY, C. & OSTDIEK, B. 2003. The Economic Value of Volatility Timing using 'Realized' Volatility. *Journal of Financial Economics,* 67**,** 473--509.

FRENCH, K. R., SCHWERT, G. W. & STAMBAUGH, R. F. 1987. Expected stock returns and volatility. *Journal of Financial Economics,* 19**,** 3--29.

FROST, P. A. & SAVARINO, J. E. 1986. An empirical Bayes approach to efficient portfolio selection. *The Journal of Financial and Quantitative Analysis,* 21**,** 293--305.

FROST, P. A. & SAVARINO, J. E. 1988. For better performance: Constraint portfolio weights. *The Journal of Portfolio Management,* 15**,** 29--34.

GARLAPPI, L. & HUANG, J. 2006. Are stocks desirable in tax-deferred accounts? *Journal of Public Economics,* 90**,** 2257--2283.

GEWEKE, J. & PORTER-HUDAK, S. 1983. The estimation and application of long memory time series models. *Journal of Time Series Analysis*, 4, 221--238.

GLOSTEN, L. R., JAGANNATHAN, R. & RUNKLE, D. E. 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance***,** 1779--1801.

GOYAL, A. & SANTA-CLARA, P. 2003. Idiosyncratic risk matters! *The Journal of Finance***,** 975--1007.

GRANGER, C. 1980. Long memory relationships and the aggregation of dynamic models. *Journal of Econometrics,* 14**,** 227--238.

GRAY, S. F. 1996. Modeling the conditional distribution of interest rates as a regime-switching process. *Journal of Financial Economics,* 42**,** 27--62.

GREEN, R. C. & HOLLIFIELD, B. 1992. When will mean-variance efficient portfolios be well diversified? *The Journal of Finance,* 47**,** 1785--1809.

HAMILTON, J. D. & SUSMEL, R. 1994. Autoregressive conditional heteroskedasticity and changes in regime. *Journal of Econometrics*, 64**,** 307--333.

HAN, Y. 2006. Asset allocation with a high dimensional latent factor stochastic volatility model. *The Review of Financial Studies*, 19**,** 237--271.

HARVEY, C. R., LIECHTY, J., LIECHTY, M. W. & MULLER, P. 2010. Portfolio selection with higher moments. *Quantitative Finance,* 10**,** 469--485.

HARVEY, C. R. & SIDDIQUE, A. 2000. Conditional skewness in asset pricing tests. *The Journal of Finance,* 55**,** 1263--1295.

HEATON, J. & LUCAS, D. 2000. Portfolio choice and asset prices: The importance of entrepreneurial risk. *The Journal of Finance,* 55**,** 1163--1198.

HUANG, J. 2008. Taxable and tax-deferred investing: a tax-arbitrage approach. *The Review of Financial Studies,* 21**,** 2173--2207.

HYDE, S., NGUYEN, N. & POON, S.-H. 2010. International Portfolio Choice and the Economic Value of Jump Timing. *Seminars Inquire Europe*. Rome, Italy.

JAGANNATHAN, R. & MA, T. 2000. Three Methods for Improving the Precision in Covariance Matrix Estimation. *Working Paper*.

JAGANNATHAN, R. & MA, T. 2003. Risk reduction in large portfolios: Why imposing the wrong constraints helps. *The Journal of Finance,* 58**,** 1651--1683.

JAMES, W. & STEIN, C. 1961. Estimation with quadratic loss. *In:* NEYMAN, J. (ed.) *The Fourth Berkeley Symposium on Mathematical Statistics and Probability*. Univ of California Press.

JEGADEESH, N. & TITMAN, S. 2001. Profitability of momentum strategies: An evaluation of alternative explanations. *The Journal of Finance*, 56**,** 699--720.

JOBSON, J. & KORKIE, B. 1981. Putting Markowitz theory to work. *The Journal of Portfolio Management*, 7**,** 70--74.

JONDEAU, E. & ROCKINGER, M. 2006. Optimal portfolio allocation under higher moments. *European Financial Management,* 12**,** 29--55.

JORION, P. 1986. Bayes-Stein estimation for portfolio analysis. *The Journal of Financial and Quantitative Analysis,* 21**,** 279--292.

JPMORGAN 1994. RiskMetrics™, Technical Document. JP Morgan.

JURCZENKO, E. F., MAILLET, B. & MERLIN, P. M. 2008. Efficient frontier for robust higher-order moment portfolio selection. Université Panthéon-Sorbonne (Paris 1), Centre d'Economie de la Sorbonne.

KANDEL, S. & STAMBAUGH, R. F. 1996. On the Predictability of Stock Returns: An Asset Allocation Perspective. *The Journal of Finance,* 51**,** 385--424.

KLEIN, R. W. & BAWA, V. S. 1976. The effect of estimation risk on optimal portfolio choice. *Journal of Financial Economics,* 3**,** 215--231.

KRAUS, A. & LITZENBERGER, R. H. 1976. Skewness preference and the valuation of risk assets. *The Journal of Finance,* 31**,** 1085--1100.

LAI, T. Y. 1991. Portfolio selection with skewness: a multiple-objective approach. *Review of Quantitative Finance and Accounting,* 1**,** 293--305.

LAMOUREUX, C. G. & LASTRAPES, W. D. 1990. Persistence in variance, structural change, and the GARCH model. *Journal of Business & Economic Statistics,* 8**,** 225--234.

LEDOIT, O. & WOLF, M. 2003. Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance,* 10**,** 603--621.

LEDOIT, O. & WOLF, M. 2004. Honey, I shrunk the sample covariance matrix. *Journal of Portfolio Management,* 30**,** 110--9.

LELAND, H. E. 2000. Optimal portfolio implementation with transaction costs and capital gains taxes. *University of California at Berkeley. Haas School of Business Working Paper*.

LETENDRE, M. A. & SMITH, G. W. 2001. Precautionary saving and portfolio allocation: DP by GMM. *Journal of Monetary Economics,* 48**,** 197--215.

LYNCH, A. W. & TAN, S. 2003. Multiple Risky Assets, Transaction Costs and Return Predictability: Implications for Portfolio Choice. *Working Paper*.

MANDELBROT, B. 1963. The variation of certain speculative prices. *The Journal of Business,* 36**,** 394--419.

MARKOWITZ, H. M. 1952. Portfolio selection. *The Journal of Finance***,** 77--91.

MARKOVITZ, H. M. 1959. *Portfolio selection: Efficient diversification of investments,* John Wiley.

MARTELLINI, L. & ZIEMANN, V. 2010. Improved Forecasts of Higher-Order Co-moments and Implications for Portfolio Selection. *The Review of Financial Studies,* 23**,** 1467--1502.

MERTON, R. C. 1971. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory,* 3**,** 373--413.

MERTON, R. C. 1973. An intertemporal capital asset pricing model. *Econometrica: Journal of the Econometric Society***,** 867--887.

MICHAUD, R. O. 1989. The Markowitz optimization enigma: is 'optimized' optimal? *Financial Analysts Journal,* 45**,** 31--42.

MODIGLIANI, F. & MODIGLIANI, L. 1997. Risk-adjusted performance. *The Journal of Portfolio Management,* 23**,** 45--54.

MOULINES, E. & SOULIER, P. 1999. Broadband log-periodogram regression of time series with long-range dependence. *Annals of Statistics*, 1415--1439.

NELSON, D. B. 1990. Stationarity and persistence in the GARCH (1, 1) model. *Econometric Theory,* 6**,** 318--334.

NELSON, D. B. 1991. Conditional heteroskedasticity in asset returns: A new approach. *Econometrica: Journal of the Econometric Society***,** 347--370.

NELSON, D. B. & CAO, C. Q. 1992. Inequality constraints in the univariate GARCH model. *Journal of Business & Economic Statistics***,** 229-235.

NIGUEZ, T. M. & RUBIA, A. 2006. Forecasting the conditional covariance matrix of a portfolio under long-run temporal dependence. *Journal of Forecasting,* 25**,** 439--458.

OWEN, J. & RABINOVITCH, R. 1983. On the class of elliptical distributions and their applications to the theory of portfolio choice. *The Journal of Finance,* 38**,** 745--752.

PÁSTOR 2000. Portfolio selection and asset pricing models. *The Journal of Finance,* 55**,** 179--223.

PESARAN, M. H. & TIMMERMANN, A. 1995. Predictability of stock returns: Robustness and economic significance. *The Journal of Finance,* 50**,** 1201--1228.

PINDYCK, R. S. 1984. Risk, inflation, and the stock market. *The American Economic Review,* 74**,** 335--351.

PLEROU, V., GOPIKRISHNAN, P., ROSENOW, B., AMARAL, L. A. N., GUHR, T. & STANLEY, H. E. 2002. Random matrix approach to cross correlations in financial data. *Physical Review E,* 65**,** 66126.

POON, S. H. 2005. *A practical guide for forecasting financial market volatility*, John Wiley & Sons Inc.

ROSS, S. A. 1976. The arbitrage theory of capital asset pricing. *Journal of Economic Theory,* 13**,** 341--360.

SAMUELSON, P. A. 1969. Lifetime portfolio selection by dynamic stochastic programming. *The Review of Economics and Statistics,* 51**,** 239--246.

SAMUELSON, P. A. 1970. The fundamental approximation theorem of portfolio analysis in terms of means, variances and higher moments. *The Review of Economic Studies,* 37**,** 537--542.

SENTANA, E. 1995. Quadratic ARCH models. *The Review of Economic Studies,* 62**,** 639--661.

SHANKEN, J. & WEINSTEIN, M. I. 2006. Economic forces and the stock market revisited. *Journal of Empirical Finance,* 13**,** 129--144.

TAYLOR, S. J. 1986. *Modelling financial time series*, Chichester, UK: John Wiley and Sons.

TEYSSIERE, G. 1998. Multivariate long-memory ARCH modelling for high frequency foreign exchange rates. Citeseer.

THORP, S. & MILUNOVICH, G. 2007. Symmetric versus asymmetric conditional covariance forecasts: Does it pay to switch? *Journal of Financial Research,* 30**,** 355--377.

TOBIN, J. 1958. Liquidity Preference as Behavior Towards Risk. *The Review of Economic Studies,* 25**,** 65--86.

VICEIRA, L. M. 2001. Optimal portfolio choice for long-horizon investors with nontradable labor income. *The Journal of Finance,* 56**,** 433--470.

VILASUSO, J. 2002. Forecasting exchange rate volatility. *Economics Letters,* 76**,** 59-64.

WEST, K., EDISON, H. & CHO, D. 1993. A Utility Based Evaluation of Some Models of Exchange Rate Variability. *Journal of International Economics,* 35**,** 23--46.

XIAO, L. & AYDEMIR, A. 2007. Volatility modelling and forecasting in finance. *In:* KNIGHT, J. & SATCHELL, S. (eds.) *Forecasting volatility in the financial markets (3 ed.).*

ZAKOIAN, J. M. 1994. Threshold heteroskedastic models. *Journal of Economic Dynamics and Control,* 18**,** 931--955.

ZUMBACH, G. 2006. The Riskmetrics 2006 Methodology. *Technical Report, RiskMetrics Group.*

ZUMBACH, G. 2009a. The Empirical Properties of Large Covariance Matrices. *RiskMetrics Journal,* 9**,** 31--54.

ZUMBACH, G. 2009b. The statistical properties of the innovations in multivariate ARCH processes with large sizes. *To appear in Quantative Finance*.